


3746 Nways Controller Models 900 and 950: IP Implementation Guide

Brian Dorling, Robert Kraus



International Technical Support Organization

<http://www.redbooks.ibm.com>



International Technical Support Organization SG24-4845-01

**3746 Nways Controller Models 900 and 950:
IP Implementation Guide**

March 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, "Special Notices" on page 351.

Second Edition (March 1998)

This edition applies to 3746 Model 900 and 3746 Nways Controller operating in a TCP/IP environment, using the 3746-9x0 IP routing feature.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1996 1998. All rights reserved.

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Preface	ix
The Team That Wrote This Redbook	ix
Comments Welcome	x
 Chapter 1. 3746 IP Router - Overview	1
1.1 Definition of Terms	2
1.2 3746 IP - Functional Overview	3
1.3 3746-900 Hardware Description	5
1.4 3746-950 Hardware Description	7
1.5 NCP-IP Routing and 3746 IP Routing	8
1.6 3746 IP to NCP Internal Connection	9
1.6.1 Internal IP Connection Definitions	10
1.7 3746 Software Requirements for IP	15
1.8 3746 Hardware Requirements for IP	16
1.8.1 IP Routing Feature	18
1.8.2 3746-950 Hardware Requirements	18
1.9 Type 3 Processors	18
1.9.1 Maximum Connectivity of the 3746-9x0 APPN/HPR Network Node	18
1.9.2 Network Node Connectivity	21
1.10 3746 Extended Functions	21
1.10.1 Activating Extended Functions	21
1.10.2 3746 Extended Functions (FC #5800)	22
1.10.3 Multiaccess Enclosure Extended Functions Part 1 (FC #5804)	23
1.10.4 Multiaccess Enclosure Extended Functions Part 2 (FC #5805)	26
1.10.5 Multiaccess Enclosure TN3270E Server (FC #5806)	28
1.11 3746 CD-ROM and Optical Disk Support	28
1.11.1 Required Action	28
1.11.2 Recommendation	29
1.11.3 Service Processor (SP) Prerequisites for CD-ROM Support	29
1.11.4 The Benefits of CD-ROM Media	29
1.11.5 Online Documentation	30
1.11.6 Stand-Alone CCM (3746 Controller Configuration and Management)	30
1.11.7 Optical Disk (OD) and CD-ROM Support Details	30
1.12 Service Processor (SP) and Network Node Processor (NNP) Support Matrix	31
1.12.1 Service Processor and Network Node Processor Feature Codes	33
1.13 Frame Relay Support History	34
1.13.1 ACF/NCP Version 6 Release 1	34
1.13.2 ACF/NCP Version 6 Release 2	34
1.13.3 ACF/NCP Version 7 Release 1	35
1.13.4 ACF/NCP Version 7 Release 2	35
1.13.5 ACF/NCP Version 7 Release 3	35
1.13.6 ACF/NCP Version 7 Release 4	36
1.13.7 ACF/NCP Version 7 Release 5	37
1.13.8 ACF/NCP Version 7 Release 6	38
1.14 3746 Microcode EC and ECA Levels	38
 Chapter 2. An Introduction to TCP/IP	43
2.1 The Internet Protocol Suite	43
2.2 An Internet	43

2.3	Client/Server Processes	44
2.4	General Description Of The TCP/IP Protocols	46
2.4.1	Transmission Control Protocol (TCP) Layer Details	50
2.4.2	Internet Protocol (IP) Layer Details	52
2.4.3	Data Link Control Layer Details	54
2.4.4	Physical Layer Details	55
2.4.5	The Socket Application Programming Interface	56
2.4.6	Application Defaults	57
2.4.7	IP Routing	58
2.5	Protocols Other Than TCP	63
2.5.1	User Datagram Protocol (UDP) Layer Details	64
2.5.2	Internet Control Message Protocol (ICMP)	65
2.6	The Domain Name System	70
2.6.1	3746 IP and the Domain Name System	71
2.7	IP Addresses	71
2.7.1	Subnetting	74
2.7.2	Broadcasting	77
2.7.3	Supernetting (Route Aggregation)	79
2.7.4	3746 IP and IP Addresses	79
2.8	Fragmentation and Reassembly	80
2.8.1	3746 IP and Fragmentation	81
2.9	Address Resolution Program (ARP)	81
2.9.1	ARP and Subnets	82
2.9.2	Proxy-ARP or Transparent Subnetting	82
2.9.3	3746 IP and ARP	84
2.10	3746-9x0 Security Considerations and Controlling IP Routing	85
2.10.1	3746 IP Filters	86
2.10.2	3746 Access Control	86
2.10.3	Rules for Filters and Access Controls	87
2.11	Bootstrap Protocol (BootP)	88
2.11.1	3746 IP and BootP	89
2.12	Request for Comments (RFCs)	90
Chapter 3.	Dynamic IP Routing Protocols - Introduction	91
3.1	Interior and Exterior Gateway Protocols	91
3.2	Choosing Gateway Protocols	92
3.3	Routing Algorithms	92
3.3.1	Static Routing	93
3.3.2	Distance Vector Routing	93
3.3.3	Link-State Routing	98
3.4	Routing Information Protocol (RIP)	99
3.4.1	Protocol Description	99
3.4.2	3746 IP and RIP	101
3.5	Open Shortest Path First (OSPF)	102
3.5.1	OSPF Terminology	102
3.5.2	Protocol Description	106
3.5.3	Multicast Extensions to OSPF (MOSPF)	115
3.5.4	3746 IP and OSPF	118
3.6	Border Gateway Protocol (BGP)	120
3.6.1	Introduction	121
3.6.2	Protocol Description	122
3.6.3	3746 IP and BGP Version 4	129
Chapter 4.	IBM 3746-9x0 and IP Routing Protocols	131

4.1 Using RIP	131
4.1.1 RIP Configuration on the IBM 3746-9x0	132
4.2 Using OSPF	134
4.2.1 OSPF Areas	134
4.2.2 OSPF Interfaces	135
4.2.3 Area Ranges	136
4.2.4 Virtual Links	136
4.2.5 OSPF Configuration on the IBM 3746-9x0	136
4.3 Using BGP Version 4	140
4.3.1 AS Numbers	141
4.3.2 BGP Configuration on the IBM 3746-9x0	141
4.4 Static Routes	145
4.5 Routing Protocols Interoperability	146
4.5.1 RIP Specifics	147
4.5.2 OSPF Specifics	148
4.5.3 BGP Specifics	149
4.5.4 Static/Default Routes	149
Chapter 5. 3746-9x0 IP Implementation Details	151
5.1 Multiple TCP/IP Stacks	151
5.2 Route Caching	153
5.2.1 Distributed Routing and ARP Caching	154
5.2.2 Default Route Caching	157
5.2.3 Displaying the CBSP Routing Tables	159
5.3 3746 IP Interfaces Overview	160
5.4 ESCON	160
5.4.1 ESCON Port Sharing	162
5.4.2 CDLC Configuration Rules	162
5.4.3 ESCON - Fiber and Host Link Configuration	165
5.4.4 ESCON - Link Stations and IP Configuration	167
5.4.5 ESCON - IP Configuration Example	168
5.5 Token-Ring	172
5.5.1 Token-Ring Port Sharing	173
5.5.2 Token-Ring - General Configuration	174
5.5.3 Token-Ring - IP Configuration	174
5.6 Ethernet	175
5.6.1 Ethernet Configuration	175
5.7 Frame Relay	176
5.7.1 Orphan Circuits	177
5.7.2 Multicast Emulation	177
5.7.3 Frame Relay Network Management	177
5.7.4 Frame Relay Scheduling	178
5.7.5 Consolidated Link Layer Management (CLLM)	184
5.7.6 Frame Relay Port Sharing	185
5.7.7 Frame Relay PVC Sharing	185
5.7.8 Frame Relay - General Configuration	190
5.7.9 Frame Relay - IP Configuration	191
5.8 Point-to-Point Protocol (PPP)	193
5.8.1 Internet Protocol Control Protocol (IPCP)	194
5.8.2 Link Control Protocol (LCP)	194
5.8.3 Network Control Protocol (NCP)	195
5.8.4 PPP BRS on the 3746 Base Adapters	195
5.8.5 PPP Basic Configuration Procedure	196
5.8.6 PPP - General Configuration	197

5.8.7 PPP - IP Configuration	197
5.9 X.25 Protocol	199
5.9.1 X.25 Port Sharing	200
5.9.2 X.25 - General Configuration	201
5.9.3 X.25 - IP Configuration	202
5.9.4 X.25 - IP Configuration Example	202
5.10 3746-9x0 IP Management	214
5.10.1 Controller Configuration and Management (CCM)	216
5.10.2 TELNET Command Line Interface	218
5.11 3746-9x0 SNMP Functions	225
5.11.1 3746-9x0 SNMP Relay	226
5.11.2 Distributed Agents	227
5.11.3 SNMP Traps	227
5.12 3746-9x0 Security	229
5.12.1 3746-9x0 IP Filters	229
5.12.2 3746-9x0 IP Access Controls	230
5.13 Changing Passwords	232
5.14 3746 IP Tuning Recommendations	235
Chapter 6. 3746-9x0 IP Operation and Configuration	237
6.1 Controller Configuration and Management (CCM)	237
6.1.1 Welcome to the CCM	237
6.1.2 Operating Environments	238
6.1.3 Installing CCM	239
6.1.4 Becoming Familiar with the CCM User Interface	244
6.1.5 About the Configuration Process	247
6.2 Using TELNET - Operation and Configuration	258
6.2.1 OPCON Command Level	258
6.2.2 GWCON Commands	262
6.2.3 CONFIG Commands	264
6.3 TELNET Line Command Examples	266
6.3.1 Using PING	267
6.3.2 Dumping the ARP Table	267
6.3.3 Dump CBSP Routing Table Display	268
6.3.4 3746-9x0 Processor IP Routing Cache Display	268
6.3.5 Specific Route Display	269
6.3.6 IP Protocol Display	269
6.3.7 Using Traceroute	270
6.3.8 Static Routes	271
6.3.9 Using Access Controls	272
6.3.10 Using Filters	275
6.3.11 Protocol Global Configuration Display	277
6.3.12 Display Frame Relay Information	278
6.3.13 Display 3746 IP Router Configuration	279
6.3.14 3746-9x0 IP Router Restart	280
Chapter 7. Using CCM - IP Configuration Examples	283
7.1 Defining the Service LAN	285
7.2 Defining Token-Ring Interface Port 2080	290
7.3 Defining Token-Ring Interface Port 2144	293
7.4 Defining Token-Ring Interface Port 2176	297
7.5 Defining Token-Ring Interface Port 2208	300
7.6 Defining ESCON Interface Port 2240	302
7.6.1 MVS Definitions	306

7.7 IP Parameter Configuration	307
7.7.1 IP General Parameter	308
7.7.2 PPP NCP Definition	309
7.7.3 IP Static Routes	310
7.7.4 IP Access Controls	310
7.7.5 IP Filters	312
7.7.6 IP BootP Forwarding	312
7.8 OSPF Parameter Configuration	313
7.9 RIP Parameter Configuration	318
7.10 BGP Parameter Configuration	321
7.11 Proxy-ARP Parameter Configuration	325
7.12 SNMP Parameter Configuration	326
Appendix A. TELNET Line Command Overview	329
A.1 OPCON Commands	329
A.2 GWCON Commands	329
A.3 Config Commands	331
A.4 IP Monitoring and Configuration	331
A.4.1 Configuring IP	332
A.4.2 Monitoring IP	333
A.5 OSPF Monitoring and Configuration	334
A.5.1 Configuring OSPF	334
A.5.2 Monitoring OSPF	335
A.6 BGP Monitoring and Configuration	336
A.6.1 Configuring BGP	336
A.6.2 Monitoring BGP	337
A.7 ARP Monitoring and Configuration	338
A.7.1 Configuring ARP	338
A.7.2 Monitoring ARP	339
A.8 BRS Monitoring and Configuration	339
A.8.1 Configuring BRS	339
A.8.2 Monitoring BRS	340
A.9 ESCON Port Monitoring and Configuration	340
A.9.1 Configuring ESCON	341
A.9.2 Monitoring ESCON	341
A.10 Token-Ring Port Monitoring and Configuration	341
A.10.1 Configuring Token-Ring	341
A.10.2 Monitoring Token-Ring	342
A.11 Frame Relay Port Monitoring and Configuration	342
A.11.1 Configuring Frame Relay	342
A.11.2 Monitoring Frame Relay	343
A.12 PPP Port Monitoring and Configuration	344
A.12.1 Configuring PPP	344
A.12.2 Monitoring PPP	344
Appendix B. 3746 IP Functions - Overview	347
B.1 TCP/IP Functions	347
B.2 Dynamic Routing Protocols	348
B.3 Interfaces Supported	348
B.4 SNMP Support	348
Appendix C. Special Notices	351
Appendix D. Related Publications	353

D.1 International Technical Support Organization Publications	353
D.2 Redbooks on CD-ROMs	353
D.3 Other Publications	353
How to Get ITSO Redbooks	355
How IBM Employees Can Get ITSO Redbooks	355
How Customers Can Get ITSO Redbooks	356
IBM Redbook Order Form	357
List of Abbreviations	359
Index	363
ITSO Redbook Evaluation	365

Preface

This redbook describes the functions, and how to configure, the stand-alone IP routing functions of the 3746-900 and the 3746-950 Nways Controllers. All 3746-9x0 IP routing functions announced July 1996 or later are discussed.

Guidelines are given to use the 3746-9x0 in an IP-only or in a hybrid IP and APPN networking environment. Several practical examples are presented to demonstrate the 3746-9x0 IP routing functions. An extensive introduction to the TCP/IP protocol suite for readers not familiar with TCP/IP is also included.

Several scenarios or examples are provided that will help the reader to understand how to customize the 3746 hardware, basic IP routing support, how to control and monitor the IP routing functions, and how to configure dynamic routing protocols running on the 3746.

A basic knowledge of networking concepts and terminology is assumed.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the Systems Management and Networking ITSO Center, Raleigh.

Brian Dorling, is an Advisory ITSO Specialist for Communications Architectures at the Systems Management and Networking ITSO Center, Raleigh. Brian is responsible for a broad range of IBM communication architectures including Advanced Peer-to-Peer Networking (APPN), Multiprotocol Transport Networking (MPTN), Networking Broadband Services (NBBS), Switched Virtual Networking (SVN), and for the 3746 family of multiprotocol communications controllers. After joining IBM in 1978, Brian has worked as a Customer Engineer and Systems Engineer in the networking field in the UK and Germany.

Robert Kraus is a TP specialist working for IBM Germany. He has three years of experience in IP implementation and debugging. He holds a degree in Informatik from the Fachhochschule Muenchen (Munich). His areas of expertise include problem determination in multiprotocol networks and designing LAN networks.

The authors of the first edition of this redbook were:

Volkert Kreuk
IBM USA

Marcus Cooke
IBM Canada

Fred Williams
IBM Canada

Thanks to the following people for their invaluable contributions to this project:

Jean-Claude Dispensa
IBM LaGaude

Gilles Garcia
IBM LaGaude

Volkert Kreuk
IBM USA

Pierre Planas
IBM LaGaude

Bruno Mansuy
IBM LaGaude

Christian Vallee
IBM LaGaude

The Editing and Graphics Team
Systems Management and Networking ITSO Center, Raleigh

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in "ITSO Redbook Evaluation" on page 365 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Web sites:

For Internet users <http://www.redbooks.ibm.com/>

For IBM Intranet users <http://w3.itso.ibm.com/>

- Send us a note at the following address:

redbook@us.ibm.com

Chapter 1. 3746 IP Router - Overview

The June 1996 networking announcements include APPN enhancements and native IP routing functions for the 3746 Model 900 and the 3746 Nways Controller. The IP routing functions make two new, high-end, IP routing-capable nodes available to IBM customers. The 3746-9x0 IP routers can be employed in an all-IP network or in a hybrid IP, SNA/APPN, and frame switching networking environment.

The introduction of the IP routing support makes the 3746-9x0 a true multiprotocol router. IP function can be run in conjunction with the NCP SNA/Subarea, APPN Network Node (NN) and Dependent Logical Unit Requester (DLUR), and frame switching functions. IP, SNA/Subarea and NN/DLUR, and frame switching functions all run natively; that is, no protocol conversion or enveloping takes place. Port and adapter sharing by these three networking architectures enable the user to build an integrated network that offer SNA/subarea, NN/DLUR, IP, and frame switching functions in a very cost-efficient way.




Figure 1. IP, APPN, and NCP/Subarea Networking. A single physical network resulting in multiple overlaying logical networks.

The 3746-9x0 is ideally suited for use in a mixed networking environment. However, it is not limited to a hybrid networking environment. The 3746-9x0 resource sharing capabilities are a very flexible user option that simplifies installation, operation and management of hybrid networks, but the 3746-9x0 is also an excellent choice for networks that run either IP, APPN, SNA/Subarea, or frame switching functions only.

It must also be noted that the 3746-9x0 is an IP router which provides scalable distributed routing of IP packets. The forwarding of IP packets is done by each processor, packets are either forwarded onto interfaces controlled by the same processor, or switched to other processors inside the 3746. The network node processor (which is required for IP) is not in any way involved in the forwarding of

IP packets. Due to the distribution of the IP forwarding across the processors, the packet forwarding capacity can be increased by adding additional processors.

The purpose of this publication is to discuss and detail the 3746-9x0 IP functions. To enable users to understand and position these functions, an introduction to the TCP/IP and dynamic IP routing protocols is also included. In this description many details on the 3746-9x0 IP implementation are included. The operation and configuration aspects of the 3746 IP router are discussed in Chapter 6, “3746-9x0 IP Operation and Configuration” on page 237, while Chapter 7, “Using CCM - IP Configuration Examples” on page 283 includes many customization examples.

IBM has announced availability of the Multiaccess Enclosure (MAE) feature for the 3746 models 900 and 950. The MAE also provides native IP routing support for the 3746. While some references are made to the MAE in this publication, this redbook focuses on the 3746 base IP routing support.

Note: *3746 Nways Controller Models 900 and 950: APPN Implementation Guide*, SG24-2536 describes, in detail, the 3746-9x0 hardware architecture and the APPN NN/DLUR functions and can be used in addition to this publication.

1.1 Definition of Terms

Within this publication the following terms will be used:

- *3746 Model 900* or *3746-900* to refer to the IBM 3746 Model 900 Expansion Unit
- *3746 Nways Controller*, *3746 Model 950* or *3746-950* to refer to the IBM 3746 Nways Multiprotocol Controller
- *3746-9x0* to refer to both 3746-900 and 3746-950, but not necessarily including the IP routing function
- *3746-900 IP* to refer to the 3746-900, including the IP routing functions
- *3746-900 NN* to refer to the 3746-900, including the APPN NN and DLUR functions
- *3746-950 IP* to refer to the 3746-950, including the IP routing functions
- *3746-950 NN* to refer to the 3746-950, including the APPN NN and DLUR functions
- *3746 IP* to refer to the IP routing functions offered by both 3746-900 IP and 3746-950 IP

The 3746 IP includes the necessary hardware and software on the 3746-9x0, the service processor and the network node processor.

- *Network node* to refer to a node providing networking function

Note: Be aware that the term network node is a general term (that is, not protocol-specific), while APPN network node refers to a node providing advanced peer-to-peer networking network node functions.

1.2 3746 IP - Functional Overview

The 3746-9X0 is a high connectivity and high throughput network node for IP, APPN/DLUR, NCP/Subarea and frame relay environments. The 3746-9X0 can control the connections of ESCON channels, token-ring LANs, Ethernet LANs, and serial (SDLC, ISDN, PPP, X.25, and frame relay) lines.

Both the 3746-900 and 3746-950 are capable of operating as an IP router. Although the 3746 Model 900 is independent of NCP for its stand-alone IP routing functions, not even requiring NCP presence, it remains connected to a 3745-XXA. The ESCON adapters of the 3746 can also be used by the NCP-IP routing function of the 3745.




Figure 2. NCP-IP and 3746 IP Adapter Support

Note: All NCP controlled functions, including the NCP IP routing function, on the 3746-900 require NCP presence. For details, refer to 1.5, “NCP-IP Routing and 3746 IP Routing” on page 8.

The 3746 Nways Controller, which is based on 3746-900 hardware and microcode, does not support the attachment of a 3745-XXA and performs its IP routing functions as a stand-alone, NCP-independent node.

The 3746-9X0 operates in conjunction with the service processor and the Network Node Processor. The Network Node Processor is a field-installable feature which will provide the 3746 network node function. As will be detailed in later chapters, the Network Node Processor provides SNMP and TELNET functions but is not involved in IP routing. The SNMP agent running on the Network Node Processor enables 3746 IP management from a remote management station, while the TELNET functions allow for remote IP operation and configuration.

The 3746 Model 900 and the 3746 Nways Controller will provide the following set of IP functions:

- IP routing providing connectivity over:

- ESCON channels

To IBM and other equipment manufacturers (OEM) processors that adhere to the ESCON architecture, possibly via one or two cascaded ESCON directors (IBM 9032 or 9033, or OEM ESCON directors).

- Token-ring
- Ethernet

Ethernet access requires one of the 3746-9x0 Ethernet features, either #5631 or #5632. Both features provide Ethernet support by physically (using a token-ring coupler) connecting the 3746-9x0 to a token-ring LAN and using an IBM 8229 translational bridge, bridging the token-ring LAN to an Ethernet LAN.

Note: Because of the translational bridging, token-ring rather than Ethernet configuration is required to enable Ethernet access to the 3746-9x0.

- Leased lines

Leased lines use either the PPP, X.25, or frame relay protocol.

- Dynamic routing protocol support, using:
 - RIP Version 1
 - OSPF Version 2
 - BGP Version 4
- ARP, InARP, and proxy-ARP
- Bootp relay

The 3746 Nways Controller provides the same high connectivity variations as the 3746 Model 900. Similar to the 3746-900, the 3746-950 allows up to sixteen adapters of any type to be configured. Using only one type of adapter, the maximum connectivity is up to one of the following:

- 16 ESCON channel couplers
- 33 token-ring couplers (32 available for customer use)
- 32 line interface couplers, with:
 - 32 high-speed lines up to 2 Mbps
 - 120 medium-speed and low-speed line interfaces up to 256 kbps

It is planned to increase this number to 240 lines in 1998.

Note: The previous text shows the maximum number of attachments supported. The number of attachments that can be simultaneously used depends on traffic volumes. Due to power constraints, the theoretical installable maximum of 960 low/medium lines (≤ 64 kbps) cannot be reached.

The distributed processor architecture of the 3746-9X0 enables high volumes of IP traffic to be routed. Its native IP over (ESCON) channel support provides the best solution currently available to enable IP access to host-resident IP applications.

IBM has carried independently audited performance tests of the 3746 compared to other OEM routers. These test were audited by the Tolly Group and are available for for viewing on the world wide web at the following URL:

http://www.tolly.com/online_1/testresults/documents/main.cfm

Users must register to access the information on this site but usage is free. Once registered, access the documents 7291, 7292, and 729 which contain the Tolly Group test results for 3746.

The 3746 IP router benefits from the 3746-9X0 architecture and technology providing high reliability, taking advantage of the backup and redundant capabilities offered, including:

- Redundant power supply with automatic non-disruptive takeover
- DC distributed power for adapters and couplers
- ESCON directors allowing backup channel connections
- Communication Line Processor automatic backup
- RSF communication capability

Additional redundancy is provided through the provision of an optional backup Network Node Processor.

The adapter hot-pluggability, the concurrent upgrade and the concurrent maintenance operations also contribute to the high availability characteristics.

1.3 3746-900 Hardware Description

As shown in Figure 3 on page 6, the 3746 Model 900 provides the IBM 3745 Communication Controllers 17A, 21A, 31A, 41A and 61A with ESCON channel adapters and additional communication line and token-ring adapters. The native support of the ESCON architecture provides flexibility in the design of host front-end installations.

The 3746 Model 900 IP router provides IP routing functions while coexisting with:

- Stand-alone APPN network and DLUR functions
- Subarea or composite network node configurations controlled by the NCP
- Frame relay frame switching functions (FRFH)

The 3746 Model 900 IP router operates independently of the APPN/DLUR and subarea function.

The maximum number of customer configurable adapter positions on the 3746-900 is 16, unless attached to a dual CCU 3745-xxA. In the latter case, the maximum number of configurable adapters is fifteen, as one of the optional adapter positions is taken by the CBTRA. For details, see Figure 4 on page 7.




Figure 3. 3746-900 Internal Structure

Each adapter contains a single processor. The ESCON processors each control one ESCON channel coupler, the token-ring processors each control two TIC3 couplers, and the communication line processors each control two or four LIC11 or LIC12 couplers. Each LIC11 supports up to 30 medium-speed lines interfaces. Each LIC12 provides a single high-speed (up to 2 Mbps) line interface. In Europe, an LIC Type 16 is offered, which provides the IBM 3746 Model 900 with a primary ISDN adapter supporting the Euro-ISDN standard. The ISDN requires the presence of ACF/NCP allowing SNA equipment to communicate with ACF/NCP over ISDN connections. IP over ISDN support is not available.

Figure 4 on page 7 depicts a 3746-900 NN attached to a dual CCU 3745-XXA. When connected to a single CCU 3745-XXA, the 3746-900 supports 16 optional processors with up to 32 coupler positions. A TIC3 coupler attaches to the CBSP2 or CBSP3, which is required to connect the 3746-900 to the service processor, and a CBC to connect to 3745-XXA CCU A.

When connected to a dual CCU 3745-XXA, the 3746-900 supports 15 optional processors, allowing a total of 31 coupler positions (15 * 2, plus the CBTRP TIC3 coupler). In addition to the CBSP2/3-controlled TIC3 for the 3746-900 to SP and the CBC to CCU-A attachment, a CBTRP (CBTRP2 or CBTRP3) is required for the CBC to CCU-B attachment. The CBTRP also supports a TIC3 coupler which is available for user traffic.




Figure 4. 3746-900 NN. Dual CCU Attachment

1.4 3746-950 Hardware Description

As depicted in Figure 5, the 3746 Nways Controller is a stand-alone APPN network node, frame relay switch, and IP router, interconnecting ESCON channels, token-rings, and serial links (SDLC, PPP, X.25, and frame relay).




Figure 5. 3746-950 Internal Structure

The ESCON channel attachment is optional, allowing the 3746-950 to be attached to a local host system. Since the 3746-950 does not require a connection to a 3745-XXA, there is no need for a Controller Bus Coupler (CBC). The spare slot on the CBSA cannot be used for another coupler and must remain empty.

For the 3746 Nways Controller similar configuration rules apply as for the 3746 Model 900. ESCON adapters, token-ring adapters, and communication line adapters can be configured in any mix.

Note: LIC16, providing ISDN access, is supported on the 3746 Model 900 only as current ISDN support requires NCP presence.

As shown in Figure 5 on page 7, the 3746-950 is a stand-alone machine not requiring an attachment to a 3745-XXA. It does require the CBSP2 or CBSP3 and a TIC3 coupler for the 3746-950 to service processor and Network Node Processor communication.

The 3746-950 has 16 optional processors and 32 available coupler slots. The second coupler position on the CBSA has been freed up by eliminating the CBC coupler. This position, however, may not be used. The 3746-950 hardware and software and the functions (most notably SNMP and TELNET) running on the network node processor together comprise a full IP stack.

1.5 NCP-IP Routing and 3746 IP Routing

In the previous sections it has been explained that on a 3746-900, NCP and 3746 IP functions can coexist. As NCP ($\geq$ V6R1) is also providing IP routing functions, it is important that a distinction is made between the 3746 IP (stand-alone) and the NCP controlled IP (NCP-IP) routing functions.

3746 IP provides IP routing for ESCON channel, token-ring (TIC3), Ethernet, and serial line (frame relay, X.25, and PPP) attached equipment. Connectivity is provided through 3746-900 attachments only. NCP-IP provides IP routing for ESCON channel, parallel channel, token-ring (TIC1/TIC2), Ethernet, and serial line (frame relay) attached equipment. However, *with the exception of ESCON*, connectivity is provided through 3745 base attachments only; that is, *not* for equipment attached to the 3746-900 (see Figure 2 on page 3).

3746 IP and NCP-IP perform their IP routing functions independently, IP connectivity between these two routers requires either an internal, or external connection. Internal (3746-3746 via the CBC) IP connections were not supported, either by NCP versions before NCP V7R6, or by the 3746 microcode before ECA 170. Refer to 1.6, "3746 IP to NCP Internal Connection" on page 9 for details on the internal 3746 IP router to NCP IP router connection support. Figure 6 on page 9 depicts how IP connectivity between NCP-IP and 3746 IP can be realized via token-ring, Ethernet, and frame relay. In all cases a port (token-ring, Ethernet, or frame relay respectively) is required on both the 3746-900 and the 3745 base frame. A single ESCON channel can also be used to connect TCP/IP for MVS ($\geq$ V3R1) to both 3746 IP and NCP-IP. However, all IP datagrams sent between 3746 IP and NCP-IP will be routed via the host router.




Figure 6. External NCP-IP to 3746 IP Routing Connections

1.6 3746 IP to NCP Internal Connection

With the introduction of NCP V7R6 it is now possible to link the NCP IP router and the 3746 IP router together via the CBC between the 3745 and 3746. Previously, an external link between the 3745 and 3746 was needed to connect the two routers. To support this function from the 3746 side, the optional feature code #5800 is needed. This is supported by 3746 microcode level ECA 170 and later levels.

The internal connection or connections between the two routers are seen as point-to-point (PtP) IP connections. These connections must be defined in the NCP and 3746 CCM. Each connection uses a separate TIC3 on the 3746 as a proxy router. This TIC3 may also be used to drive a token-ring, but beaconing on that token-ring or other problems may interfere with the IP routing function. For that reason it is recommended, but not mandatory, that each TIC3 used for an internal IP connection is used only for that function.

Figure 7 on page 10 shows a representation of the internal point-to-point IP connection.




Figure 7. Internal IP PtP Connection

1.6.1 Internal IP Connection Definitions

The following sections explain how to define an internal point-to-point IP connection. Figure 8 on page 11 shows the configuration that is being defined. The definitions shown here include all the information needed for this part of the NCP generation process. The actual PtP link is defined in the LN2240 LINE **1** and IP2240 PU **2** statements and the IPLOCAL statement for LADDR=10.04.00.99 **5**. The NETWORK must be IP **3**, and PUTYPE must be 1 **4** (see Figure 9 on page 12 and Figure 10 on page 13).




Figure 8. Example Configuration

1.6.1.1 NCP Definitions

Figure 9 on page 12 and Figure 10 on page 13 show the NCP definitions needed to configure the internal IP PtP link from the 3745 side.

```

*****
* VTAM HOST DEFINITIONS
*****
*
HOSTA01  HOST  BFRPAD=0,DELAY=0.0,INBFRS=6,MAXBFRU=32,
              SUBAREA=1,
              STATMOD=YES,UNITSZ=1024
*
*****
*   TOKEN RING LINE 2240
*****
*
ODLCGRP1 GROUP  ECLTYPE=(PHY,ANY),ADAPTER=TIC3,
              DIAL=NO,LNCTL=SDLC,TYPE=NCP
LN2240   LINE   ADDRESS=(2240,FULL) 1,
              LOCADD=400000002240,
              MAXPU=2,
              MAXTSL=16732,
              NPACOLL=YES,
              PORTADD=2,
              SPEED=9600,
              TRSPEED=16
IP2240   PU     ADDR=02,
              INTFACE=(TR2240,2048) 2,
              NPACOLL=YES,
              NETWORK=IP 3,
              PUTYPE=1 4,
              ARPTAB=(1)

*****
*                               CHANNEL ADAPTERS
*****
*
CAGRP1   GROUP  LNCTL=CA,CA=TYPE7,NCPCA=ACTIVE,MONLINK=CONT,
              TIMEOUT=240.0,ISTATUS=ACTIVE,CASDL=420.0,DELAY=0.0
*
*****
* VTAM HOST CHANNEL ATTACHMENT
*****
*
CALNP1   LINE   ADDRESS=P1
CAPUP1   PU     PUTYPE=5,TGN=1
*
*****
* CHANNELS FOR IP CHANNEL ATTACHED ROUTERS
*****
*
CALNP2   LINE   ADDRESS=P2,CASDL=420,DELAY=0.0,MONLINK=NO
CAPUP2   PU     PUTYPE=1,INTFACE=CPU2,ARPTAB=(10,,NOTCANON)

```

Figure 9. NCP Definitions


```

*****
*           IP ROUTING DEFINITIONS
*****
*
      IPOWNER HOSTADDR=21.1.1.1,NUMROUTE=(25,25,25),UDPPORT=580,
              INTERFACE=(CPU2),MAXHELLO=6,NUMDRIF=10

      IPLOCAL LADDR=21.1.1.99,INTERFACE=CPU2,METRIC=1,
              P2PDEST=21.1.1.1,PROTOCOL=RIP
*
      IPLOCAL LADDR=10.04.00.99 5, INTERFACE=TR2240, METRIC=1,
              P2PDEST=10.04.00.01 6, PROTOCOL=RIP, SNETMASK=255.255.0.0
*
      IPROUTE DESTADDR=13.0.0.0 7,
              NEXTADDR=10.4.00.01 8,
              INTERFACE=TR2240,
              METRIC=1,DISP=PERM,HOSTRT=NO
*
      GENEND      GENEND
                  END

```

Figure 10. NCP Definitions

For reference, the following is a short description of some of the parameters used:

- IPOWNER** Identifies the TCP/IP MVS/VM host that manages the routing tables (NCPRROUTE) for the NCP-IP router.
- IPLOCAL** Defines an interface to the NCP IP router.
- IPROUTE** Defines an entry in the NCP IP routing table.
- HOSTADDR** Defines the IP address of the owning IP host.
- INTERFACE** Defines the name of the IP interface to NCPRROUTE.
- LADDR** Defines the IP address of the interface.
- P2PDEST** Defines IP address of the destination IP host.
- SNETMASK** Defines subnet mask for the interface.

1.6.1.2 3746 CCM Definitions

The following section describes how to use CCM to create the 3746 definitions for the PtP link.

1. The token-ring port 2880 must be configured first, and IP must be activated on that port.




Figure 11. Port Configuration

2. On port 2880, an IP address must be defined. This is the IP address of the 3746 end of the PtP IP connection (10.4.0.1). This must match the IP address specified on the P2PDEST operand of the IPLOCAL **6** statement in NCP.




Figure 12. IP over Token-Ring Parameters

3. RIP should be activated as shown on port 2880.




Figure 13. OSPF/RIP Parameters per IP Address

4. The following RIP parameters should be selected for IP address 10.4.0.1.




Figure 14. RIP Parameters per IP Address

To use the internal PtP link for NCP-IP routing to reach the destination network 13.0.0.0 (DESTADDR=13.0.0.0 **7**) from NCP, the next hop (NEXTADDR **8**) points to the IP address of the 3746 end of the PtP connection (10.4.0.1).

1.7 3746 Software Requirements for IP

Software requirements for using the 3746 IP function are:

- TCP/IP for MVS V3R1 (+PTFs) and higher for ESCON attachments.
TCP/IP for MVS V3R2 has improved performance and is recommended.
- OS/390 TCP/IP stack V2R4 and higher.
- ACF/NCP V6R1 for NCP-IP to 3746 IP connectivity over Ethernet.
- ACF/NCP V7R1 for NCP-IP to 3746 IP connectivity over token-ring.
- ACF/NCP V7R3 for NCP-IP to 3746 IP connectivity over frame relay.
- ACF/NCP V7R6 for NCP-IP to 3746 IP connectivity over the internal PtP link.

- The IBM 3746 IP router is supported by Router and Bridge Manager (RABM), RABM is presently available with Nways Campus Manager LAN which is shipped as a component of Nways Manager for AIX V1R2 (order number is 4300291). RABM V2 and alert transport also ship as part of Nways Enterprise Manager.

Distributed Console Access Facility (DCAF) V1R3, or TME Remote Control V1R0, are required to provide remote access from a customer's controlled station to the service processor and the Network Node Processor.

Note: DCAF can provide remote console support over either SNA or TCP/IP.

1.8 3746 Hardware Requirements for IP

The 3746 IP function on 3746-900s shipped after June 1995 requires the following:

- The network node base (#9018 or #9023) feature

The network node base feature provides the required hardware for the 3746-900. It includes the new Controller Bus and Service Processor (CBSP) Type 2 or 3.

- The network node processor (NNP) (#5022 or #5122) feature

The network node control feature consists of the required hardware and software for the 3746-900 Network Node Processor; it includes the processor hardware, and the IP router configuration and management software.

A second network node processor feature may be ordered to increase Network Node Processor point availability. Note, that in an IP environment the Network Node Processor is responsible for SNMP and TELNET functions only, while the IP forwarding functions and dynamic routing protocols operate in the 3746-9x0 processors. An inactive Network Node Processor therefore, will impact TELNET and SNMP functions but does not affect the IP forwarding functions.

Note: A prerequisite for the network node processor feature is the service processor (#5021 or #5052). A service processor (#5020) configured with the field installable feature service processor upgrade (#5026) can also be used.

- The controller expansion (#5023) feature

The controller expansion (#5023) is a corequisite feature required to house the Network Node Processor and service processor. This 19-inch rack assembly may be bolted to the side of the 3746-900.

Note: Up to two controller expansions can be ordered. Only a single rack can be mounted to the 3746-9x0.

- Optionally, the side cover (#5024) feature

A side cover (#5024) feature needs to be ordered if the controller expansion (#5023) is not bolted to the side of the 3746-900.

- Optionally, the Ethernet feature (#5631 or #5632) to enable Ethernet-attached equipment to access the 3746 IP functions

The Ethernet port feature (#5631) provides a TIC3 coupler together with an Ethernet - token-ring bridge (IBM 8229) and an IBM 8228 token-ring medium access unit (MAU) integrated in the 3746-9x0. The integrated Ethernet - TR bridge feature (#5632) provides just the token-ring to Ethernet bridge and the MAU.




Figure 15. 3746-9x0 IP Ethernet Support

Note: A similar configuration can be obtained with customer-installed equipment. However, when ordering the Ethernet feature, IBM will be responsible for correct installation and configuration.

The 3746 IP function on 3746-900s shipped before July 1995 requires the same features as for a new 3746-900, except that instead of the network node base (#9018) feature, the network node base upgrade (#5123) feature is required. The NN base upgrade feature consists of a CBSP-to-CBSP3 upgrade. In addition, the service processor upgrade (#5026) must be ordered for installations still using the service processor Type 1 (#5020, or #5021 type 9585); this upgrade feature is not required when using the service processor Type 2 (#5021), or a #5052.

It should be emphasized that 3746 IP functions require attachment to the more powerful Type 2 or Type 3 ESCON or token-ring processors. Type 2 processors provide greater processing power and twice the memory of a Type 1 processor, Type 3 processors provide greater processing power and twice the memory of a Type 2 processor. Subarea functions on the 3746-900 will continue to be supported on the Type 1 and 2 processors. The new communication line processors (CLP3), support subarea, APPN NN, and IP functions.

Note: The 3746-950 is equipped with Type 2 or Type 3 processors only.

Important

The upgrade of Type 1 to Type 3 processors, if required, is a responsibility of the customer and is *not* included in the 3746-9x0 features discussed.

Summarizing the relevant features to upgrade a 3746-900 to enable IP support:

- The network node upgrade (#5019 or #5123) feature
- The network node processor (#5022 or #5122) feature
- Optionally, the service processor upgrade (#5026) feature

The service processor feature #5052 is recommended for customers presently using the #5020 feature.

- The controller expansion (#5023) feature
- Optionally, the side cover (#5024) feature
- Optionally, ESCP, TRP, and CBTRP upgrade to ESCP3, TRP3, and CBTRP3 respectively

1.8.1 IP Routing Feature

The microcode to run the 3746 IP functions is included in the network node processor feature (#5022 or #5122). However, to enable the 3746 IP routing support, the IP routing feature has to be ordered (IP routing feature, #5033) separately for each 3746-9x0 on which you want to run IP. The IP routing feature provides a software key that is unique to the 3746-9x0 it is ordered for.

1.8.2 3746-950 Hardware Requirements

The 3746 Nways Controller does not support a 3745-XXA attachment and operates as a truly stand-alone node. 3746-950 IP functions on the 3746-950 require the same features as for a new 3746-900.

Summarizing:

- The network node processor (#5022 or #5122) feature
- The controller expansion (#5023) feature
- Optionally, the side cover (#5024) feature

The network node processor feature (#5022 or #5122) is mandatory when ordering the 3746 Nways Controller. The 3746-950 can operate in both IP and APPN networking environments. 3746-950 customization and operation are fully independent of NCP and VTAM. The 3746-900 is field upgradeable to the 3746-950.

1.9 Type 3 Processors

In September 1997 new *type 3* processors became available for the 3746. These processors provide increased performance and double the memory size of type 2 processors. They increase APPN/DLUR performance by up to 100%, and connectivity by up to 200% over type 2 processors.

For 3746 machines running IP, the CBSP3 provides improved dynamic routing protocol performance and provides greater storage for routing tables. The other type 3 processors also provide more space for routing table and ARP caches plus better packet forwarding performance.

Field upgrades are available from type 1 and type 2 processors to type 3 processors.

Table 1 on page 19 shows individual processor connectivity and Table 2 on page 21 shows total 3746 connectivity.

1.9.1 Maximum Connectivity of the 3746-9x0 APPN/HPR Network Node

The number of PUs, frame relay DLCIs, and sessions available on the 3746-9x0 are given in the following tables.

1.9.1.1 Adapter Connectivity

Table 1 gives the maximum number of PUs, frame relay DLCIs, and APPN or dependent LU sessions that the various 3746-9x0 adapters can handle, assuming that the IP routing software is not loaded in these adapters.

For adapters providing IP routing, the maximum number of PUs and sessions controlled by the 3746 network node may be lower, due to the storage used by the IP routing software.

Depending on the storage available in the processors, the actual maximum number of 3746-controlled PUs and sessions may be different. The maximum number of ESCON logical link stations (16) and, in case of 3746 Model 900, the maximum number of NCP-controlled PUs (see column NCP in Table 1) and total number of PUs (see column Total in Table 1) are absolute maximum numbers which cannot be exceeded.

<i>Table 1. Adapter Level Connectivity</i>						
Adapter	3746 Model 900				3746 Model 950	
	PUs ¹			Sessions ² 3746 NN	PUs ¹	Sessions ²
	3746 NN	NCP	Total			
ESCP	0	16	16	0	-	-
ESCP2	16 ⁹	16	16 ⁹	4900	16 ⁹	4900
ESCP3	16 ⁹	16	16 ⁹	14 000	16 ⁹	14 000
TRP	0	2000	2000	0	-	-
TRP2¹⁰	1000	2000	2000	4500	1000	4500
TRP3¹⁰	2000	2000	2000	13 500	2000	13 500
For CCU B³:						
TRP	0	500	500	0	-	-
TRP2¹⁰	1000	2000	2000	4000	-	-
TRP3¹⁰	2000	2000	2000	13 000	-	-
CBSP	-	500	500	-	-	-
CBSP2/CBSP3⁴:	-	500	500	-	-	-
CBSP2/CBSP3⁵:	0	0	0	0	0	0
CLP with:						
3000 DLCIs⁴	-	4000 ⁶	4000 ⁶	-	-	-
500 DLCIs¹⁰	1000 ⁸	2000 ⁷	2000 ⁷	3500	1000 ⁸	3500
CLP with:						
3000 DLCIs⁴	-	4000 ⁶	4000 ⁶	-	-	-
2000 DLCIs¹⁰	3000 ¹²	3000 ¹¹	3000 ¹¹	12 500	3000 ¹²	12 500
Legend: CBSP2 Controller bus and service processor (type 2) CBSP3 Controller bus and service processor (type 3) CCU Central control unit CLP Communication line processor CLP3 Communication line processor (type 3) DLCI Data link connection identifier ESCP2 ESCON processor (type 2) ESCP3 ESCON processor (type 3) LU Logical unit NN Network node PU Physical unit TRP2 Token-ring processor (type 2) TRP3 Token-ring processor (type 3)						

The following notes refer to Table 1 on page 19.

Notes:

1. These are adjacent PUs (or ESCON logical link stations), such as end nodes, network nodes, LEN nodes, dependent PUs, gateway downstream PUs, and X.25 virtual circuits. For the 3746-900, the total of NCP-controlled and 3746-controlled stations can not exceed the total that is in the Total column.
2. These are all the LU sessions (independent and dependent LUs) routed by the 3746 adapter, including LU-LU sessions involving non-adjacent nodes. HPR/ANR sessions between HPR/RTP nodes that do not begin or end in the 3746 are not part of these numbers and can be any number. For the 3746-900, these numbers do not include the sessions routed by NCP. The quantity of NCP-routed sessions depends on the 3745 storage capacity.

These figures apply only to processors that have a few PUs or ESCON stations.
3. This is the TRP, TRP2, or TRP3 used to connect the 3745 CCU-B to the 3746-900.
4. For a 3746-900, if neither 3746 APPN/HPR nor 3746 IP routing is used in any CLP/CLP3.
5. For any 3746-950, and any 3746-900 using the 3746 APPN/HPR network node or IP routing support.
6. Up to 1000 SDLC PUs and any mix of up to 3000 frame relay PUs, ISDN PUs, and X.25 virtual circuits (one PVC or SVC per PU).
7. Up to 1000 SDLC PUs and any mix of up to 1000 frame relay PUs, ISDN PUs, and X.25 virtual circuits (one PU per PVC or SVC).
8. Up to 1000 PUs over SDLC, frame relay, and X.25 lines.
9. This includes any logical stations (TCP/IP) used by the 3746 IP Router.
10. Not all the maximum connection capabilities are possible simultaneously. For a given processor, the maximum number of resources in a category (3746-controlled PUs, NCP-controlled PUs, 3746-controlled sessions, SDLC links) depends on the number of active resources in other categories, on the presence of the IP routing feature, and, in case of a CLP, on the mix of lines (SDLC, frame relay and X.25).

For example: TRP2s (without the IP routing feature) support simultaneously a total of 500 APPN/HPR PUs and 3000 data sessions, or 1000 dependent PUs and 1500 data sessions.
11. Up to 1000 SDLC PUs and any mix of up to 2000 frame relay PUs, ISDN PUs, and X.25 virtual circuits (one PU per PVC or SVC).
12. Up to 1000 SDLC PUs and any mix of up to 2000 frame relay PUs and X.25 virtual circuits (one PU per PVC or SVC).

1.9.2 Network Node Connectivity

Table 2 gives the total number of PUs, APPN and dependent LU sessions, and lines that a 3746 network node can handle (no matter what (type 2) adapter configuration is used).

Table 2. Network Node Level Connectivity		
Connectivity		Comments
Type	Number	
PU	5000	End nodes, LEN nodes, network nodes, Dependent PUs.
Sessions	15000	All the LU-LU sessions using 3746 DLUR and APPN routing, including sessions involving non-adjacent nodes. HPR/ANR sessions between HPR nodes connected to the 3746 are in addition to this number of sessions and can be in any quantity.
CLP or Serial Lines	120	Frame relay, SDLC, X.25, and PPP.
Note: For the 3746 Model 900, the resources beyond these network node quantities are controlled by NCP(s) either as part of a PU type 4 (SNA) node or part of an APPN composite network node (CNN).		

1.10 3746 Extended Functions

From microcode level D46130I (ECA 170) onwards a new system of password usage to activate chargeable features will be implemented in the 3746. This system has the following objectives:

- Allows the user to access new functions included in one or more feature codes, each controlled by a password.
- The functions are associated to with each 3746-9x0 and are controlled by different passwords on each machine.
- The customer pays for the functions used on each 3746-9x0.

1.10.1 Activating Extended Functions

Figure 16 on page 22 shows how extended functions are activated, and where the passwords are entered.




Figure 16. Activating Extended Functions

The following sections give an explanation of how functions have been divided into the new feature codes.

1.10.2 3746 Extended Functions (FC #5800)

Feature code #5800 must be ordered to get the corresponding password and operate any of the following functions.

1.10.2.1 Functions Controlled by the NNP (3746 Models 900 and 950)

The following functions are controlled by the NNP:

- HPR MLTG support over token-ring, Ethernet, SDLC, frame relay, and X.25
- Bandwidth Reservation System (BRS) for IP traffic over PPP lines
- Frame relay enhancements:
 - Frame relay switching (FRFH) - CIR
 - Bandwidth Reservation System (at DLCI level) between IP, APPN, and HPR
- X.25 (requires FC #5030):
 - SNA/DLUR, APPN, HPR, and IP over X.25
 - X.25 port sharing between NNP and NCP traffic (not NPSI)
 - PVCs and SVCs
 - NPM support (requires NPM V2R3 + PTFs)

1.10.2.2 Functions Controlled by NCP (3746 Model 900)

The following functions are controlled by NCP:

- Internal IP coupling to 3746 IP router (requires NCP V7R6)
- Dynamic windowing enhancements for frame relay/ISDN (requires NCP 7.6)
- PRI Euro-ISDN (LIC16/FC.5216) enhancements (requires NCP V7R5+):
 - Automatic backup of frame relay links over ISDN (non-disruptive for NNP to NCP connections)
 - NPM support (requires NPM V2R3 + PTFs)

1.10.3 Multiaccess Enclosure Extended Functions Part 1 (FC #5804)

Feature code #5804 must be ordered in order to get the corresponding password and operate any of the following functions supported via the multiaccess enclosure of the 3746 Models 900 and 950:

- Interactive Network Dispatcher (NetDispatcher)
- RIP V2 (on multiaccess enclosure ports only)
- Native HPR over ATM
- Enhanced ATM adapters: LIC294, LIC295
- Branch Extender
- Enhancements announced in September 1997:
 - ATM
 - Frame relay, PPP, ISDN and WAN
 - APPN/HPR, DLSw and IP

1.10.3.1 Detailed Description of FC #5804

The MAE Extended Functions provide a set of multiprotocol routing protocols and transport software to enable scalability and load-balancing capabilities for S/390 IP/Web servers connected to the Internet or intranet:

- Interactive Network Dispatcher

The Network Dispatcher function provides load balancing among a set of IP servers adjacent to the router running this function. The load-balancing mechanism uses technology from IBM's Research Division to determine the most appropriate server to receive each new connection. Subsequent traffic for that connection is then forwarded to the same server. The routing is transparent to users and other applications. The load information is obtained from a set of weights based upon the number of connections active per server, the number of new connections since the last interval, feedback from response time of individual HTTP, FTP, SSL servers, and configurable policy information.

The Network Dispatcher sees only the incoming packets from the client to the server. It does not need to see the outgoing packets, which significantly reduces the overhead imposed by load balancing. The client's packet is forwarded to the chosen server exactly as it was created. Since Network Dispatcher is also available on AIX, Windows NT, and Sun Solaris, it is useful for many applications such as e-mail servers, Web servers, distributed parallel database queries and other TCP/IP applications. The Interactive Network Dispatcher feature for Multiaccess Enclosure is priced based on the number of servers being balanced.

Included with the Multiaccess Enclosure Extended functions is a license for Network Dispatcher for IBM Networking (5801-AAR FC #2151). One Network Dispatcher for IBM Networking Use Authorization license (program 5807-AAR FC #1453) must be purchased for each server. The user may not exceed the number of servers for which they are authorized.

- RIP Version 2

RIP Version 2 adds the following features: route tags to propagate EGP information, subnet masks to support variable subnet masks, next hop addresses to support optimization of routes, authentication for password

passing, and multicasting so that multicast can be used instead of broadcast. RIP V2 is available today on multiaccess enclosure (FC #3000) ports only.

- ATM:
 - High-Performance ATM Adapters
1-port 155 Mbps MMF (FC #3294 - LIC294) and 1-port 155 Mbps SMF (feature number 3295 - LIC295) provide improved performance compared to LIC284 and LIC293 respectively.
 - Native APPN/HPR over ATM
APPN/HPR supports native ATM so that the router can attach directly to ATM network without LAN emulation or encapsulation. This support includes: ATM signaling of bandwidth, QoS, ATM addressing, connection network support for SVCs, route selection extensions for ATM characteristics, mapping between ARB and ATM characteristics, and HPR over ATM MIB extensions.
 - Native ATM bridging allows routers to connect frame relay/ATM interworking switches to devices on either PVCs or SVCs that do not support LAN emulation connections.
 - Configurable Quality of Service (QoS) allows LAN emulation networks to take advantage of ATM's QoS capabilities.
 - Next Hop Resolution Protocol (NHRP) enables shortcut routes for IP across ATM networks. NHRP supports zero-hop routing for endstations with NHRP and one-hop routing for stations without NHRP clients.
 - For added failure recovery, a backup gateway for endstations on LAN emulation can now be configured with default gateway IP addresses. If the primary gateway goes down, the backup gateway automatically starts passing packets from the endstation to other subnets. Additionally, the user can configure which ARP server is the primary and backup.
 - Server Cache Synchronization Protocol (SCSP) distributes the SRP servers to eliminate a single point of failure.
 - All the supported routed protocols and native ATM bridging may be multiplexed onto a single ATM permanent virtual circuit.
- Branch Extender
Branch Extender is an APPN option that can be used to build a large APPN/HPR network that delivers access to customers' existing SNA applications with scalability and cost-effectiveness. It provides a gateway function to interconnect thousands of branch offices to an enterprise wide area network (WAN) and improves efficiency of network flows.

The network node with Branch Extender addresses the problem of too many network nodes by limiting the gateway node to the topology of the branches it serves on its downstream links and uses a form of default routing for network services on its upstream link. In this way, Branch Extender reduces topology database size and traffic, adds the ability to register resources from a branch to a central directory server in the WAN, and allows for direct connections between subnetworks without using an extended border node. This function is available for any data link type supported by APPN.
- Frame relay, PPP, ISDN, and WAN enhancements:

- Frame relay

Frame relay dial circuit interface is configurable on a V.25bis interface type.

Frame relay data compression (mode-1 FRF.9) is configurable per PVC to run over a frame relay interface.

Congestion can be reduced with support for congestion management via CLLM messages: SNMP traps are sent on receipt of CLLM, FECN, or BECN frames, the 3746 reduces the transmission rate upon receipt of a FECN, and sends a BECN.

- PPP

PPP Bandwidth Allocation Protocol/Bandwidth Allocation Control Protocol (BAP/BACP) adds the ability to dynamically add or drop links over ISDN B channels.

Authentication servers can now be used so that names and passwords need not be configured at each router.

Encryption Control Protocol (ECP) using Data Encryption Standard (DES) Cipher Block Chaining (CBC) mode is now available for PPP.

- ISDN

ISDN I.430 and I.431 is supported to enable interconnecting to the leased-line service from NTT.

- Ethernet

The Ethernet locally administered MAC address can be configured to override the default burned-in address.

- WAN

Bandwidth Reservation (BRS) can assign TCP/IP packets to a BRS class and priority based on the packet's UDP or TCP port number.

A backup frame relay, PPP, or X.25 link can be specified for IP over frame relay when the traffic rate reaches a specified threshold.

Enabling or disabling of adapters can be done from a single operator console without knowing which interface(s) is configured for WAN reroute.

- APPN and DLSw enhancements:

- APPN/HPR

Native HPR over ATM (see ATM section)

Implicit focal point and up to eight backups enable the router to initiate a management session with NetView.

- DLSw

A range of source/destination SAPs and MAC addresses can be configured to override circuit priority.

The DLSw Switch-to-Switch Protocol allows the exchange of a MAC address list between partners.

NetBIOS session alive spoofing eliminates session alive frames on a dial-on-demand link.

1.10.4 Multiaccess Enclosure Extended Functions Part 2 (FC #5805)

FC #5804 must be ordered to get the corresponding password. FC #5805 is required to operate any of the following functions supported via the Multiaccess Enclosure (MAE) of the 3746 Models 900 and 950:

- FDDI, HSSI, F-Enet, 128 MB (MAE memory)
- ESCON MultiPath Channel (MPC+) support for IP applications
- Enterprise Extender for APPN/HPR over IP backbone
- Channelized T1/E1 and other enhancements

1.10.4.1 Detailed Description of FC #5805

FC #5805 requires FC #5804 and provides the following additional functions for the MAE (FC #3000 and #3001):

- FDDI support
One FDDI interface per adapter (FC #3286 - LIC286). Operates as either a Dual Attach Station (DAS) or a Single Attach Station (SAS) using multimode fiber (MMF).
- HSSI support
One HSSI interface per adapter (FC #3289 - LIC289). Supports T3 and E3 speeds.
- 10/100 Mbps Ethernet support
One interface per adapter (FC #3288 - LIC288) with speeds of either 10 Mbps or 100 Mbps.
- 64 MB Additional system memory (FC #3520)
An optional second 64 MB DIMM is available on Multiaccess Enclosure for a total of 128 MB of memory for especially demanding environments (TN3270E, DLSw, APPN, DLUR, etc.).
- High-Performance Data Transport (HPDT) for UDP
High-Performance Data Transfer (HPDT) MultiPath Channel (MPC), also known as MPC+, has been extended to include IP support over the ESCON channel. HPDT UDP extends the efficiencies of HPDT services to applications using the OS/390 UNIX System Services UDP interface. HPDT reduces CPU cycle consumption and achieves a more efficient transfer of data. HPDT UDP is initially targeted for communications between DB2 on OS/390 V2R4 (requires PTF PQ04890) and SAP R/3 application servers. Other UNIX System Services socket applications using UDP, such as NFS and DCE, can also transparently take advantage of HPDT UDP services over the Multiaccess Enclosure ESCON channel.
- High-Performance Data Transport (HPDT) for TCP
HPDT TCP/IP extends the efficiencies of HPDT services to IP applications using OS/390 V2R5. HPDT reduces CPU cycle consumption and achieves a more efficient transfer of data. It is supported over the ESCON channel of the MAE.
- Enterprise Extender
Enterprise Extender is a simple set of extensions to APPN High-Performance Routing (HPR) technology to integrate SNA into IP backbones. To the HPR

network, the IP backbone is a logical link; to the IP network, the SNA traffic is UDP datagrams.

Enterprise Extender provides the flexibility for SNA parallel sysplex features that are currently available in HPR networks now to be available to users in networks that have IP backbones, or even IP clients when coupled with TN3270e server support. Enterprise Extender also makes it possible for SNA networks to use IP attachments as alternate and backup routes for the SNA network.

Enterprise Extender technology can also reduce the demands on the data center routing platforms, and, thus, provide a more cost-effective solution than other integration technologies. Enterprise Extender seamlessly routes packets through the network protocol edges eliminating the need to perform costly protocol translation and the store-and-forward associated with transport layer functions.

The Enterprise Extender technology also provides many of the traffic control features that SNA users have come to expect. Using Class of Service (COS), SNA applications specify the nature of the services they require from the network (for example, interactive, batch, etc.). Enterprise Extender supports SNA priority in IP environments by mapping the SNA COS priority to UDP port numbers that can be easily prioritized using Bandwidth Reservation System (BRS).

- Channellized T1/E1 support

This support allows the current ISDN PRI adapter to be configured as a channellized T1 or E1 in lieu of using it for ISDN PRI. Support is provided for frame relay and PPP over individual or groups of DS0s. One or multiple connections are supported on the same physical interface. The bandwidth of each connection will be a multiple of 64 kbps up to the maximum speed of 24*64 for T1 or 31*64 for E1. The time slots for the combined DS0s need not be contiguous.

- Dial-in support for SDLC PU Type 2 devices

Switched dial-in is the capability for SDLC PU Type 2 devices to dial through a switched data network. The support will be provided through DLSw. It will provide HDX and FDX support as well as NRZ and NRZI. Call answering is supported but a dial-out facility is not.

- X.25 scalability on the Multiaccess Enclosure

This extends the current X.25 capacity from a limit of 239 VCs to a limit that is memory-dependent and capable of supporting more than 1000 VCs.

- RIP outage-only advertisements:

Allows RIP advertisements to only be sent on an interface when a route is missing from the IP route table. This will facilitate advertisement on ISDN/V.25bis Dial-on-Demand (DoD) circuits only in circumstances where the DoD circuit also has data traffic to send. Circuits will not be brought up solely for RIP advertisements.

1.10.5 Multiaccess Enclosure TN3270E Server (FC #5806)

This feature must be ordered to get the corresponding password and operate the TN3270E server. This function is supported via the Multiaccess Enclosure of the 3746 Models 900 and 950 and requires the MAE Extended Function FC #5804.

The TN3270E server provides a TN3270 gateway function for TN3270 clients downstream from an SNA/VTAM S/390. The clients connect to the server using a TCP connection which the server then maps to a corresponding SNA LU-LU session that the TN3270 server maintains with the S/390.

The TN3270E server supports the capabilities defined in RFCs 1576, 1646 and 1647.

The connectivity from the TN3270E server to the S/390 uses DLUR in the Multiaccess Enclosure and is supported, locally and remotely, over all the interfaces that support DLUR. When coupled with Enterprise Extender, the TN3270 servers can be distributed in the network with an IP infrastructure and, therefore, be placed in locations that provide the best scalability and availability without regard to backbone protocol.

1.11 3746 CD-ROM and Optical Disk Support

3746 Licensed Internal Code supporting new functions available December 1997 (or later) will be delivered on CD-ROM only (media FC #9300). Customers with an optical disk reader on their service processor, who wish to use new functions *must* upgrade their service processor or install the latest service processor.

For service processors not equipped with the CD-ROM reader, the Optical Disk media remains available after December 1997 for microcode maintenance up to ECA 170/EC level D46130x. (x = alphabetic letter identifying the technical level of the microcode; x includes possible future levels.) It will also be available for shipment of 3746 MES which does not require any of the new functions available December 1997 or later. For example, an order for processors Type 3 (ESCP3/TRP3/CLP3/CBSP3) is in this category and may include the Optical Disk media (FC #9500).

1.11.1 Required Action

Configurations not yet submitted to AAS must be updated before they are entered in AAS. If you enter an order based on an old CFREPORT file, and your order includes functions available December 1997 or later, make sure that the media feature included in the configuration is FC #9300 (CD-ROM).

Important Reminder

Any order without a service processor *must* include a media feature number (either #9300 or #9500). This requirement applies whether this order is for a new 3746 or for a 3746 MES. (This is applicable to all shipments on or after December 1997.)

1.11.2 Recommendation

CF3745 was enhanced November 1997 to automatically determine the media and service processor features required to support the hardware/microcode functions being configured for a 3746 machine or MES.

It is strongly recommended that existing orders/configurations are validated via CF3745 before entering new orders or altering orders planned to be shipped December 1997 or later.

1.11.3 Service Processor (SP) Prerequisites for CD-ROM Support

The following is a reminder of the SP support for the 3746 functions available starting December 1997:

- Service Processor Type 2 (feature 5052): 3746 functions are supported.
- Service Processors Type 1 (feature 5021): 3746 functions are supported if the SP is equipped with the 64 MB memory expansion (feature 5028) and a 2 GB hard disk drive.

The 3172-based service processors (P/N 41H7520 and P/N 55H7630) are equipped with a 2 GB hard disk drive, but the 9585-based service processor requires the SP hard disk drive Upgrade (feature 5026)

- Oldest Service Processors (feature 5020): 3746 functions are not supported. A Service Processor Type 2 (feature 5052) must be ordered.

Note: The Service Processor Type 2 contains a CD-ROM reader. For the new functions available December 1997 or later, or to benefit from CD-ROM specific functionalities (see below), the Service Processors Type 1 (feature 5021) must be upgraded with the previously mentioned features and with the SP CD-ROM Upgrade (feature 5051). However, it is strongly recommended not to upgrade with these features but to go directly to the Service Processor Type 2 (FC #5052).

1.11.4 The Benefits of CD-ROM Media

Using CD-ROM media provides the following benefits:

Dual Load Module Support

Two levels of 3746 microcode can be installed on the hard disk drive (hard disk drive) of the Service Processor. Benefits of this new support are discussed in the following sections.

Online Microcode Upgrades

The 3746 remains operational during the installation of a new microcode level. While the new microcode is being installed as the non-active level, the 3746 remains in production on the current active level. If the Service Processor (SP) operating system is modified, the SP functions such as CCM support, are not available during this step.

The 3746 is interrupted only during activation of the non-active level (re-IML), which lasts about 5 to 15 minutes, depending on the number of processors in the 3746 configuration. The customer may perform this activation at a later time, when compatible with the network operations. Compared to a microcode upgrade via optical disk media, the 3746 interruption time is reduced from 150 minutes (average) to less than 15 minutes.

Backup Production Microcode Level

After the new microcode level is activated, the previous production level becomes non-active, but remains available for possible re-activation if for any reason the customer needs to switch back to this production level. Should this happen, the 3746 interruption time is again limited to the re-IML time of the 3746.

Testing the Non-Active Microcode Level

The dual level of code also allows the users to test a new level of microcode, maintenance level or functional level, by taking advantage of unused machine time; they can activate a trial level of microcode and then come back to the production level when required.

1.11.5 Online Documentation

The CD-ROM media containing the new Licensed Internal Code of 12/12/97 will also include the following 3746 product information:

- All the customer publications, including the *3745/3746 Overview*, GA33-0180-8, *3745/3746 Planning Guide*, GA33-0457-0, *CCM User's Guide* and service publications (installation and maintenance, etc.).
- Most of the 3746 presentations and documents available today on the IBM intranet home page (<http://w3.lagaude.ibm.com/ccp/3746.htm>).

Netscape and Acrobat Reader (included in the CD-ROM) allow the user to display the above information at the service processor screen.

1.11.6 Stand-Alone CCM (3746 Controller Configuration and Management)

The following possibility is under study for the near future: all supported levels of the stand-alone CCM would be available on the CD-ROM containing the Licensed Internal Code. This would allow the user to retrieve the CCM level corresponding to the microcode level of his/her 3746(s) and run this CCM on a workstation. Also under study is a simple procedure (no need for diskettes) for installing this CCM on the workstation.

1.11.7 Optical Disk (OD) and CD-ROM Support Details

Data migration from OD to CD-ROM is supported from any microcode level supporting optical disk to any microcode level supporting CD-ROM. Customer parameters (MOSS-E) are saved on diskette and no longer on a second optical disk.

The following are the microcode EC levels:

- D46130x

D46130 is the *last* microcode EC supporting OD media (and the previous microcode installation methodology). See Table 3 on page 31 for service processor requirements.

- F12380x - December 1997 GA level of microcode

F12380 is the *first* microcode EC supporting CD-ROM media and only this media (and the new microcode installation methodology). See Table 3 on page 31 for required service processor features.

Table 3. Service Processor Requirements		
Service Processor Type	Microcode Delivery Media	
	CD-ROM	Optical disk
Microcode Level	F12380X	up to D46130X
7585 (Pentium-200Mhz) FC 5052	Yes 1 2 With CD enabled	Yes 2 With OD enabled
3172 (Pentium-90Mhz) FC 5021	Yes With FC 5051 With FC 5028	Yes 3
3172 (486-66Mhz) FC 5021	Yes 4 With FC 5051 With FC 5028	Yes 3 4
9585 (486-66Mhz) FC 5021	Yes 4 With FC 5051 With FC 5028 With FC 5026	Yes 3 4
9577 (486-33Mhz) FC 5020	No	Yes 3 4

CD-ROM operations require the service processor to be equipped with a CD-ROM drive, 2-GB hard disk drive and 96 MB of memory.

Notes:

1 Starting December 1997, SP Type 7585 is equipped with a CD-ROM reader only and no OD drive, 96-MB memory and 2-GB hard disk drive.

2 Between June 1997 and December 1997, SP Type 7585 was equipped with an OD drive and a CD-ROM reader. The OD drive was enabled as default. Starting December 1997, CD-ROM media will be used for shipment of microcode when an MES is ordered for a 3746 using such an SP (CD-ROM media, FC.9300, included in the MES by CF3745). The existing OD drive will be removed and the CD-ROM drive enabled when the MES is installed.

3 See requirements for FC 5026 and FC 5028, depending on functions configured on the 3746 (APPN, MAE, etc.). Refer to CF3745 for detailed information.

4 For better response time at the operator console, it is strongly recommended that this SP is replaced by a 7585 (FC 5052).

1.12 Service Processor (SP) and Network Node Processor (NNP) Support Matrix

The following table gives an overview of the service processor (SP) and network node processor (NNP) features and machine types needed to support each of the available 3746 hardware configurations. All numbers prefixed with a # are feature codes.

Table 4 (Page 1 of 2). Service Processor (SP) / Network Node Processor (NNP)
Feature Code and Machine Type

Function	Service Processor					NNP	
	#5020	#5021			#5052 1	#5022	#5122 2
	9577 3	9585 4	3172 5	3172 6	7585 7	3172 8	7585 9
NNP traffic only	Yes	Yes	Yes	Yes	Yes	11	11
Second expansion enclosure (#5016)	10	#5028	#5028	#5028	Yes	11	11
SP Sharing	10	#5028	#5028	#5028	Yes	11	11
APPN/HPR (NNP) IP (#5033)	#5026	Yes	Yes	Yes	Yes	Yes	Yes
APPN/ISR/DLUR/RTP More than: -3000 PUs -9000 LU-LU sessions	#5026	Yes	Yes	Yes	Yes	#5027	#5027
APPN/ISR/DLUR/RTP More than 15000 LUs-LUs sessions (up to 30000) 12	#5052	#5026 #5028 #5051 10	#5028 #5051 10	#5028 #5051	Yes	#5211	Yes
3746 -#5802 (SSE) -More than 120 lines controlled by the NNP (up to 240) -CD-ROM support -APING from SP	#5052	#5026 #5028 #5051 10	#5028 #5051 10	#5028 #5051	Yes	Yes	Yes
MAE -#3000/#3001 -Extended function (#5804) -Extended functions 2 (#5805) -TN3270e server (#5806) -HSSI (T3/E3 speeds) (#3289) -Fast Ethernet (#3288) -FDDI (#3286) -64MB memory expansion (#3520)	#5052	#5026 #5028 #5051 10	#5028 #5051 10	#5028 #5051	Yes	Yes	Yes

Table 4 (Page 2 of 2). Service Processor (SP) / Network Node Processor (NNP)
Feature Code and Machine Type

Function	Service Processor					NNP	
	#5020	#5021			#5052 1	#5022	#5122 2
	9577 3	9585 4	3172 5	3172 6	7585 7	3172 8	7585 9
Previews EBN	#5052	#5026 #5028 #5051 10	#5028 #5051 10	#5028 #5051	Yes	Yes	Yes
Notes: 1 Service Processor Type 2 2 Network Node Processor Type 2 3 486-33Mhz 4 486-66Mhz 5 486-66Mhz P/N 41H7520 6 Pentium-90Mhz P/N 55H7630 7 Pentium-200Mhz 8 Pentium-90Mhz 9 Pentium-200Mhz 10 Recommended alternative: Replace the service processor by a service processor type 2 (#5052) for better response times at the operator console. 11 No requirement on the NNP. 12 Requires also a CBSP3 or CBSP3 upgrade.							

1.12.1 Service Processor and Network Node Processor Feature Codes

The following is a description of SP and NNP feature codes.

- #5020** Service Processor (Type: 9577)
- #5021** Service Processor (Type: 9585, 3172 P/N 41H7520 or 3172 P/N 55H7630)
- #5022** Network Node Processor (Type: 3172 P/N 41H7522, no longer orderable)
- #5026** Service Processor HDD upgrade (HDD 2GB/1GB formatted)
- #5027** Network Node Processor Memory Expansion (64MB available on #5022 and #5122)
- #5028** Service Processor Memory Expansion (64 MB available on #5021)
- #5029** Service Processor Rack Mount Kit (for SP re-installation in the Controller Expansion #5023). #5029 may be required if rack mount equipment is not available. It contains:
 - 2 brackets support and 1 plate for the Display Screen
 - 2 brackets support for any kind of Service Processor (but 7585)
 - 1 drawer kit for Keyboard/Mouse
 - 1 plate for the Modem and Optical Disk Drive
- #5051** Service Processor CD-ROM drive (Provides a CD-ROM drive for SP #5021)

- #5052 Service Processor Type 2 (Type: 7585 with 96MB memory, 2GB HDD, Cd-ROM drive)
- #5122 Network Node Processor Type 2 (Type: 7585)

1.13 Frame Relay Support History

In addition to being an APPN/SNA node or an IP router, the 3746 is also a frame relay switching node. For IP and APPN/SNA traffic, it also supports connection to a frame relay network as a frame relay terminating equipment (FRTE) node. The following sections give an overview of the history of frame relay support in ACF/NCP, 3745, and 3746.

1.13.1 ACF/NCP Version 6 Release 1

NCP provided its first frame relay support in NCP V6R1, which was available August 1992.

NCP V6R1 runs on all non-A models of the 3745 Communications Controller family, providing a wide range of capacities and throughput. This frame relay support is implemented strictly in the NCP software without requiring any new hardware. It runs on the existing 3745 hardware and line adapters, including the T1/E1 High Speed Scanner. No host application software changes are required to use NCP frame relay.

Functionally, NCP V6R1 provides the ability to interconnect NCPs to each other through a frame relay network, providing NCP-to-NCP subarea traffic support. NCP acts as a frame relay endstation, called Frame Relay Terminal Equipment (FRTE). This is referred to as NCP's INN FRTE function.

Note: A frame relay endstation is known as an FRTE but is sometimes referred to as an FR Data Terminal Equipment (DTE) analogous with X.25 terminology.

1.13.2 ACF/NCP Version 6 Release 2

NCP V6R2 (June 1993) provided a general purpose frame relay switching capability for multiprotocol, SNA and non-SNA, traffic. The switching capability is referred to as the Frame Relay Frame Handler (FRFH) function.

This allows connection of FRTEs to a network of 3745s and provides end-to-end connections (PVCs) for FRTE pairs. NCP V6R2 runs on all models of the 3745 Communications Controller and on existing line adapters, including the High-Speed Scanner (HSS). NCP frame relay allows attachment using leased lines at speeds from 600 bits per second up to 2 Mbps. FRFH and INN FRTE functions can be shared on a single frame relay line.

Note: The FRFH function is often referred to as FR Data Communications Equipment in analogous (DCE) with X.25 terminology. In NCP publications the term Frame Relay Switching Equipment (FRSE) is used as well.

In addition, LMI support was added in this release.

1.13.3 ACF/NCP Version 7 Release 1

NCP V7R1 (January 1994) provides SNA peripheral device connectivity via frame relay. This is referred to as NCP's BNN FRTE function, and uses the RFC1490 routed frame format. BNN FRTE, FRFH and INN FRTE functions can be shared on a single frame relay line. IBM networking equipment such as the IBM AS/400, IBM 3174, IBM 3172 and RouteXpander/2 can be used to directly access the SNA network, making use of the boundary function, and also transfer other frame relay traffic over different virtual circuits using the NCP FRFH support.

The frame relay BNN capability of ACF/NCP V7R1 supports multiple stations per DLCI for 3745 adapters (not 3746-900 attached).

1.13.4 ACF/NCP Version 7 Release 2

NCP V7R2 (October 1994) provides support for frame relay connections on the 3746 Model 900 that use the LIC11 and LIC12 adapters. The frame relay support on the 3746 Model 900 is similar to the support in the base frame. This support includes:

FRTE support

- Subarea frame relay connections to other NCPs.
- Frame relay virtual circuits from peripheral devices to the boundary function in NCP.

FRFH support

Frame relay switching between lines on the 3746 Model 900 may be defined in NCP. The information is passed to the Communication Line Adapters (CLAs) on the 3746 Model 900, which provide outboard frame switching totally independent of the 3745.

3746 LMI support

LMI support was added for 3746 Model 900 connected lines.

NCP V7R2 introduces communications rate (CR) support, which allows users to allocate a minimum bandwidth to each virtual circuit, depending on the traffic needs of the corresponding endstations. This guarantees that traffic will flow on a given virtual circuit at least at its communications rate. At the same time, any unused bandwidth is made available for use by other active virtual circuits.

Note: The communication rate must be defined per PVC segment. An end-to-end communication rate requires consistent definitions on each of the PVC segments that comprise a virtual circuit.

1.13.5 ACF/NCP Version 7 Release 3

NCP V7R3 (March, 1995) had the following frame relay support enhancement:

Multiple BNN stations per DLCI for 3746 lines

Support of multiple BNN stations per DLCI on 3746 connected lines was added.

Frame relay over token-ring

Frame relay over token-ring gives ACF/NCP Version 7 Release 3 the capability to support frame relay frame handler functions (FRFHs) between NCPs over token-ring (IEEE 802.5) physical connections. This will allow customers who interconnect NCPs with token-ring to provide a private frame relay network over these token-ring connections.

Frame relay over token-ring requires ACF/VTAM Version 4 Release 2 with the appropriate PTF (supported on 3745 base adapters only).

IP over frame relay

IP over frame relay is an enhancement to NCP's frame relay function that allows IP frames to be transmitted over and received from a frame relay network without being encapsulated in SNA frames. This is termed native IP routing. Previously IP traffic routed from one 3745/NCP to another 3745/NCP over a communications link required SNA encapsulation.

With the implementation of IP over frame relay, NCP Version 7 Release 3 also supports dynamic reconfiguration (DR) of IP frame relay interfaces. It is possible to permanently define a frame relay IP PU on a frame relay physical link that does not already have IP resources defined. This requires the use of the permanent dynamic reconfiguration function available with VTAM Version 4 Release 3.

The IP over frame relay capability is RFC 1490 compliant.

Frame relay BAN support

The frame relay boundary access node (BAN), which uses the RFC149 bridged frame format, provides an extension to the previously announced frame relay boundary network node (BNN) capability. BAN supports APPN and BNN traffic on 3745 and 3746 connected lines.

Increased DLCI range

The DLCI range is increased from 16-215 to 16-991 for all lines. This allows greater flexibility in specifying DLCI numbers to match those assigned by frame relay providers when attaching to a frame relay network.

DLCI sharing

NCP V7R3 supports the following DLCI sharing options:

- One INN station and IP per DLCI (3745 only)
- One INN station per DLCI (3746 only)
- Multiple BNN/APPN stations and IP (3745 only)
- Multiple BNN/APPN stations per DLCI (3746 only)

1.13.6 ACF/NCP Version 7 Release 4

NCP V7R4 (March 1996) had the following frame relay support enhancements:

Frame relay internal frame switching support

This function will provide customers the ability to couple the 3745/NCP frame relay frame handler support function (FRFH) with the 3746 connected lines. For those customers using an external frame relay line to switch traffic between 3745 and 3746, this function will allow them to eliminate this external line by defining an internal PVC segment between a 3745 base line (either a frame relay physical line or NTRI frame handler logical line) and a 3746 line so that traffic can be switched internally. This allows users to frame switch between a 3746 Model 900 frame relay line and either a 3745 frame relay line or a 3745 Token-Ring Interface Coupler (TIC1 or TIC2) supporting frame relay over token-ring.

Improvements in performance, usability, and the elimination of the cost of external connections are advantages that can be achieved by frame relay users who employ this capability.

1.13.7 ACF/NCP Version 7 Release 5

NCP V7R5 (November 1996) had the following frame relay support enhancements:

Frame relay port sharing

NCP V7R5 together with 3746 microcode level D46120B allows 3746 frame relay ports to be shared by NCP, 3746 NNP and 3746 IP functions.

Boundary access node (BAN) for subarea links

Boundary access node (bridged format) support over subarea (PU4) connections is provided as an additional frame relay connectivity option. This support allows 3746 Model 900 connections to remote token-ring capable physical units (PUs) through a remote BAN router such as an IBM 6611 or IBM 2210. The 3745 does not support BAN for subarea traffic.

Frame relay inoperative error count management upgrades

Previously, if NCP received 64 consecutive frames in error, the frame relay physical link and all its associated resources were brought down. This change allows the user to configure the allowed number of consecutive error seconds before the physical link resources are brought down. This allows a frame relay physical line to survive a large burst of errors occurring in a very short time. With this change, NCP only counts one error in each 100 ms period for which an error occurs thus allowing greater network control for the system manager.

3746 DLCI sharing

3746 microcode level D46120B together with NCP V7R5 supports the following 3746 DLCI sharing options:

- One INN station (routed or bridged frame format) per DLCI
- Multiple BNN/APPN stations (routed or bridged frame format) controlled by NCP or NNP, and 3746 IP per DLCI

Enhanced dynamic windowing algorithm

NCP uses the IEEE 802.2 dynamic window algorithm in place of a Committed Information Rate (CIR) implementation to regulate the amount of bandwidth each logical SNA station is offered. NCP V7R5 provides two changes to support an enhanced dynamic windowing algorithm for 3745 frame relay logical lines:

1. To more closely approximate a CIR for permanent virtual circuits (PVCs) that are assigned a low CIR on a high-speed link, NCP enables the transmission rate for a logical SNA station to be slowed even further than the dynamic windowing algorithm allows. NCP V7R5 introduces a variable time delay between I-frame transmissions when network congestion continues while the station's dynamic working window is 1. In order to keep the delays within reason, NCP V7R5 adds a new configuration parameter to specify the upper boundary of this delay.
2. To prevent NCP from over-reacting to mild congestion, NCP ignores subsequent Backward Explicit Congestion Notification (BECN) for

100ms after an initial BECN is received and the dynamic working window is adjusted. NCP V7R5 adds a new configuration parameter to adjust this timer.

1.13.8 ACF/NCP Version 7 Release 6

NCP V7R6 (September 1997) had the following frame relay support enhancements:

NCP communications rate enhancements

NCP COMRATE enhancements allow the definition of a relative priority for each protocol that is sharing a DLCI (FRRATES keyword).

NCP dynamic windowing enhancements

The granularity of the enhanced dynamic windowing algorithm frame loss (Dw) and congestion detection (Dwc) suboperands has been improved.

Switched frame relay lines

NCP controlled switched frame relay lines carrying SNA peripheral and subarea traffic are now supported.

Congestion control parameters set via NDF

3746 attached frame relay lines can now have their congestion control parameters (for FECN, BECN and DE) set via NDF parameters. Previously this could only be set from MOSS-E.

1.14 3746 Microcode EC and ECA Levels

The following describes the functional levels of microcode (ECA), starting with the latest announced ones (possibly not yet available). For each level, only the newly supported functions are indicated, as they are cumulative with the functions of all the other levels down in the list (older levels). Once a new level is available, all the new machines (and hardware MES that ship with microcode) are automatically shipped at this level of microcode. The new level is not distributed to the installed 3746-900s and 3746-950s, except some levels such as the engineering change (EC) level D22510K (ECA number 142) of 12/95.

If the minimum EC level for a requested function is not installed or planned to be installed (for example, as part of an MES) order the most current ECA number via the IBM service representative. When an MES is shipped with microcode (see description below), this microcode is shipped at the most current functional level (ECA number) and technical level (EC number).

The EC number indicated for a given ECA number reflects the technical level of the microcode at the availability date of the ECA. The alphabetic index (A, B, etc.) complementing the EC number is incremented when a new technical level starts shipping. Starting with ECA 167, the operator can display the active ECA and EC numbers of the 3746-9x0(s) connected to the service processor:

1. EC level F12380x (ECA number 175)

Availability: December 1997

Automatically shipped with MES containing one of the following feature numbers: 3294, 3295, 5033, 5052, 5122 and 5804. If the minimum EC level D46130D is not installed, the new level is shipped with any MES containing feature numbers 3000, 3287, 5051, 5123, 5203, 5523, 5623, and 5800. If minimum EC level D46120 is not installed, the new level is shipped with any

MES containing one of the following feature numbers: 5016, 5027, 5028, 5030, 5216, 5631 or 5032.

This ECA may be ordered only for *installed* 3746s running at a lower ECA number/EC level and requiring one of the new microcode functions listed hereafter:

- APING from service processor to APPN/HPR PUs
- Dual level of code (active and non-active)
- CD-ROM support (reduced configuration/maintenance time)

For features 5052 (SP Type 2) and 5051 (SP Upgrade Type 1) shipped after December 1997 this microcode is automatically shipped.

2. EC level D46130I (ECA number 170)

Availability: 30/11/97 (older levels no longer orderable)

Shipped on OPTICAL DISK only (for SP with Optical Disk Drive) For service processors configured for CD-ROM support, see ECA 175.

If not installed, automatically shipped with the first MES containing one of the following feature numbers: 3287, 5016, 5022, 5027, 5028, 5030, 5033, 5122, 5123, 5203, 5216, 5523, 5623, 5631, 5632, 5800

This ECA may be ordered only for INSTALLED 3746s running at a lower ECA number/EC level and requiring microcode function(s) which are not orderable via feature code, such as the functional improvements of older ECAs (167 and before).

In particular, feature number 5800 (3746 Extended Functions 1) must be ordered to operate any of the following functions:

- 3746-NCP internal IP coupling (requires NCP V7R6)
- Dynamic windowing enhancements for frame relay/ISDN (NCP 7.6)
- X.25 traffic controlled by the NNP (independent from NCP)
- 3746-controlled traffic:
 - SNA/DLUR, APPN, HPR and IP
 - PVCs and SVCs
 - X.25 port: sharing between 3746 NN, 3746 IP and NCP ODLC traffic
 - NPM support (NPM V2R3 + PTFs)
- HPR MLTG support over TR, Ethernet, SDLC, frame relay and X.25
- Bandwidth Reservation System (BRS): 3746 PPP lines (IP traffic)
- Frame relay switching (FRFH): 3746 lines controlled by the NNP
- CIR for 3746 frame relay lines controlled by the NNP
- Bandwidth Reservation System at DLCI level: IP, APPN, HPR
- PRI Euro-ISDN (LIC16/FC.5216) enhancements (3746-900/NCP V7R5+):
 - Automatic backup of frame relay links over ISDN (nondisruptive for NCP to NCP connections)
 - Support of LIC16 for NPM (NPM V2R3 + PTFs)

3. EC level D46130B (ECA number 167)

Availability: June 1997

Automatically shipped with MES containing one of the following feature numbers: 3000, 5022, 5033, 5052 or 5122. If the minimum EC level D46120 is not installed, the new level is shipped with any MES containing one of the following feature numbers: 5016, 5027, 5028, 5030, 5216, 5631 or 5032.

This ECA may be ordered only for *installed* 3746s running at a lower ECA number/EC level and requiring one of the new microcode functions listed hereafter:

- HPR/RTP and ARB: TR, Ethernet, SDLC, frame relay and ESCON
- 3746 NN connectivity increase: More than 4000 PUs (up to 5000) and more than 12000 APPN/DLUR data sessions (up to 15000)
- CCM: Configuration checking versus installed configuration (CDF-E)
- Display of the 3746 microcode level (ECA, EC) at MOSS-E console

4. EC level D46120A (ECA number 159)

Availability: December 1996

- IP over leased lines (PPP) and frame relay lines
- NPM support: Ports/stations (APPN/HPR), Processor/TIC3 utilization
- Frame relay line and DLCI sharing: NCP (V7R5), 3746 NN and 3746 IP

5. EC level D46120 (ECA number 157)

Availability: November 1996

- 3746 NN connectivity increase: 4000 PUs + 12000 APPN/DLUR data sessions per 3746 NN.
- Second expansion enclosure (FC.5016): Six processors (total 16).
- Network node processor memory expansion (FC.5027).
- Service processor memory expansion (FC.5028).
- EMEA only: LIC16 (FC.5216) for Euro.ISDN primary support by NCP V7R5 (3746-900).
- FR BAN for INN traffic with remote 372x connected to 2210/2216 (NCP V7R5).
- Backup network node processor (FC.5022).
- Optimization of processor storage utilization: Increased APPN and DLUR connectivity (more PUs/sessions per processor). Up to 100% connectivity increase.
- Selective load of microcode (APPN/HPR, IP) by category of processors: ESCP2, CLP, TRP2 and (CB)TRP2.
- NNP no longer required for DCAF TCP/IP consoles on service LANs.

6. EC level D46100A (ECA number 155)

Availability: September 1996

- Ethernet interfaces (FC.5631, FC.5632): APPN/DLUR, HPR, IP, etc.
- 3746 IP Routing (FC.5033) over ESCON, token-ring, Ethernet
- 3746 HPR/ANR over ESCON, TR, Ethernet, SDLC, frame relay

- 3746 APPN/DLUR over frame relay (BNN and BAN)
- Maximum connectivity:
 - 3746-950: 500 PUs per CLP
 - 3746-900: 1000 PUs per CLP, including NCP-controlled frame relay PUs and X.25 PUs
- APPN/DLUR performance improvement by 30-40% in transactions/sec and data throughput
- Port sharing:
 - ESCON port: traffic for 3746 NN, 3746 IP, NCP-A and NCP-B
 - TIC3 port: traffic for 3746 NN, 3746 IP and one NCP
- NetView Topology Manager (3746 APPN/HPR): 3746 local topology
- Network Management (3746 IP): SNMP (NetView/AIX), Telnet, CCM
- CCM: 3746 online configuration changes, including ESCON, with:
 - Automatic deactivation/activation of impacted resources
 - Delete/Copy/Search functions
 - 3746-900 with NNP: ESCON Generation Assistant replaced by CCM
- Year 2000 support (applicable also to 3745 Models xxA)
- Network node processor installation option: No loading of the APPN/HPR microcode in the 3746-900 processors

7. EC level D22560D (ECA number 146)

Availability: 03/29/96

- 3746 NN support: LIC11, DLUR for LIC11/LIC12, four LICs per CLP, CLP backup, 16 versus four host link stations per ESCP2, 3000 versus 1500 LU sessions per TRP2, 2000 versus 1000 adjacent nodes per 3746 NN and 6000 versus 3000 LU-LU sessions per 3746 NN
- X.25 Support - FC 5030 (3746-900/NCP V7R4)
- Non-ERP support over TIC3 for CNN ANR (3746-900/NCP V7R4)
- 3746-900: 3000 versus 1000 station (aggregate number of PUs over frame relay and virtual circuits over X.25)

8. EC level D22560A (ECA number 144)

Availability: January 1996

- 3746 NN support (3746-900, 3746-950): APPN over LIC11, TIC3, and ESCON. DLUR over TIC3. Network node processor (FC.5022).
- EGA integrated to MOSS-E (3746-900/NCP).
- 2000 versus 500 PUs on (CB)TRP2 (3746/NCP).

9. EC level D22510K (ECA number 142)

Availability: December 1995 - automatically shipped to all the installed 3746-900s.

- V.24 interface (LIC11) support at speeds up to 28.8 kbps (V.35 modems)
- Native IP routing over ESCON channels (3746-900/NCP V7R3)

- Frame relay frame switching between 3745 & 3746-900 (NCP V7R4)
- ISDN Terminal Adapter (BRI/PRI) attachment (3746-900/V7R2)
- CLP: 1000 stations (aggregate number of PUs over frame relay and virtual circuits over X.25), in addition to 1000 SDLC PUs

10. EC level D22510J (ECA number 138)

Availability: August 1995

- CLP: Frame relay BAN and frame relay SAP multiplexing (3746-900/NCP V7R3)

11. EC level D22510I (ECA number 137)

Availability: June 1995

- CLP: X.25 (3746-900/NCP V7R3/NPSI V3R8)
- NPM 6-900/NCP V7R3)
- ISDN Terminal Adapter (BRI/PRI) attachment (3746-900/V7R2)
- CLP: 1000 stations (aggregate number of PUs over frame relay and virtual circuits over X.25), in addition to 1000 SDLC PUs

12. EC level D22510D (ECA number 134)

Availability: September 1994

- CLP: Frame relay (3746-900/NCP V7R2)
- CLP: Up to 120 versus 100 active lines
- CLP: Up to 1,000 versus 500 active PUs

Chapter 2. An Introduction to TCP/IP

TCP/IP is a set of networking protocols developed to allow cooperating computers to share resources across a network. TCP/IP products are made by different vendors, and tens of thousands of networks of all kinds use them.

2.1 The Internet Protocol Suite

TCP/IP is a set of data communication protocols that are referred to as the *internet protocol (IP) suite*. Because TCP and IP are the best known of the protocols, it has become common to use the term TCP/IP to refer to the whole family. TCP and IP are two of the protocols in this suite. Other protocols that are part of the Internet suite are UDP, ICMP, and ARP. Each of the protocols is discussed in the following sections.

2.2 An Internet

An Internet is a collection of *heterogeneous* networks using TCP/IP. The administrative responsibilities for an Internet (for example, to assign IP addresses and domain names) can be within a single group or distributed among multiple groups. Networks comprising an internetwork can use either the same or different technologies. For example (see Figure 17), within network A the attached stations are using local area network (LAN) technology (token-ring, Ethernet, FDDI, etc.) to communicate, while within network B communication is made possible through wide area network (WAN) technology (frame relay, PPP connections, X.25, etc.).




Figure 17. An Internet

Host stations can send messages from any of them to any other, except where there are security or other policy restrictions on access. Communication within a single network is referred to as *intranetworking*, and communication between stations that are attached to different networks is called *internetworking*. Stations within a single network can communicate directly.

Note: Within TCP/IP the term *network* is used to refer to the logical entity that enables direct station-to-station communication (also called a *subnet*, see 2.7.1, "Subnetting" on page 74), but also used as a synonym for an Internet. When referring to an Internet, this publication will use the term *IP network*.

Internetworking requires the presence of one or multiple internetworking devices (for example, device R). Internetworking devices are referred to as IP routers, as they *route* data from one network into the other. No principal distinction exists between a host station and a router. On both devices the same IP protocol suite is required. The only difference is that a router is attached to multiple networks and provides routing functions between these. In older publications, IP routers are often referred to as *gateways*. In this publication we will use the more common term *router*. The 3746 IP router is an example of an IP router or IP gateway.

2.3 Client/Server Processes

TCP/IP is a set of protocols that provide *lower-layer* functions needed for TCP/IP services. Examples of TCP/IP services are:

- File Transfer Program (FTP)

FTP allows a user on any computer to get files from another computer, or to send files to another computer. File transfer requires that a *session* be started first between a local (client) process and a remote (server) process (TCP is used as the transport layer). A session is started by specifying a station to connect to. Commands to transmit or receive files are issued to the local FTP process and forwarded to the remote process.




Figure 18. File Transfer - FTP

FTP client and server functions are not supported on the 3746 IP router. However, it will perform IP routing functions between two stations employed in file transfer.

- Remote Login (TELNET)

The network terminal protocol (TELNET) allows a user to log in on any other computer on the network. TELNET requires that a session be started first between a local (client) process and a remote (server) process (TCP is used as the transport layer). A session is started by specifying a station to connect to.




Figure 19. Remote Login - TELNET

The TELNET program effectively makes the local (client) station invisible. Anything entered on the local TELNET process is forwarded to the remote processor. Responses are returned and displayed on the local station. Implementations of TELNET generally include a terminal emulator for some common type of terminal. Most common are VT100 and VT200 terminal emulation.

The 3746 IP has implemented a TELNET server that enables remote access to the 3746 IP operator functions. For details on the 3746 IP TELNET functions, refer to 5.10, “3746-9x0 IP Management” on page 214.

- Network Management (SNMP)

The simple network management protocol (SNMP) provides access to management data maintained on any other computer on the network. Using a remote (server) process, SNMP allows a local (client) station to retrieve data from the remote computer or set variables on the remote computer.




Figure 20. Network Management - SNMP

The data retrieval can be *solicited*, that is, data is returned after the client process has explicitly requested it, or it can be *unsolicited*. Unsolicited data is sent to the client process when an EVENT (for example, link up or link down) occurs. Unsolicited data is referred to as *traps*. The SNMP GET and GET NEXT commands are used to retrieve solicited data. The SNMP SET command is used to set variables on the remote station. The SNMP TRAP command allows the remote station to send unsolicited data.

SNMP does not require a session to be started before data can be exchanged (SNMP uses UDP as a transport layer). Solicited data is sent in response to a

query. Unsolicited data is sent immediately, but requires that the possible recipients of traps are preconfigured on the server (that is, agent) station.

SNMP management is based on a manager-agent relation. A network management station (for example, TME 10 NetView for AIX) provides SNMP client functions that are used to retrieve and set management information maintained on SNMP agents (SNMP servers) running on remote stations. To make the information exchange between manager and agents implementation independent, SNMP data is defined and referenced using ISO's Abstract Syntax Notation 1 (ASN.1). ASN.1 is a formal language defined by the ISO that removes any possible ambiguity from both representation and meaning.

Configuration and performance data that is maintained on the SNMP agents is referred to using format descriptions in so-called *management information bases (MIBs)*. Proper management support requires that the management stations are aware of all MIBs supported on the agent station. The most common MIB is MIB-II, which describes the network interfaces, data sent and received, etc. Frame relay, ESCON, token-ring, and many other MIBs exist to reference to more specific data maintained on the agents.

The 3746 IP router supports SNMP agents functions, which enable remote SNMP management from a network management station. For details on the 3746 IP SNMP functions, refer to 5.11, "3746-9x0 SNMP Functions" on page 225.

As can be seen in the description of the previous examples, application services provided within the framework of TCP/IP all use a *server/client* model. A server function provides a specific service for the rest of the network, while a client function is used to access the service. For more details about the application aspects of TCP/IP, see 2.4.5, "The Socket Application Programming Interface" on page 56.

Note: Server and client need not be on different computers. They could be different programs running on the same computer.

2.4 General Description Of The TCP/IP Protocols

TCP/IP is a layered set of protocols. In order to understand what this means, it is useful to look at an example. A typical situation is the file transfer protocol. First, there is the FTP protocol. This defines a set of commands that one machine sends to another (that is, commands to specify a user ID and password), what file to retrieve, and in which directory to store the file. FTP, like other application protocols, simply defines a set of commands and messages to be sent but, using TCP and IP, assumes that there is a reliable way to communicate.

TCP is responsible for making sure that the commands get through to the other end. It keeps track of what is sent and retransmits anything that did not get through. If a file is too large for one datagram, TCP will split it up into several datagrams and make sure that they all arrive correctly. Since these functions are needed for many applications, they are put together into a separate protocol, rather than being part of the specifications for file transfer.

You can think of TCP as forming a library of routines that applications can use when they need reliable network communications with another computer. Similarly, TCP calls on the services of IP. Although the services that TCP supplies are needed by many applications, there are still some kinds of applications that do not

need them. However, there are some services that every application needs. These services are put together into IP.

As with TCP, you can think of IP as a library of routines that TCP calls on, but which is also available to applications that do not use TCP. The strategy of building several levels of protocol is called *layering*.




Figure 21. TCP/IP - A Layered Networking Structure

Applications programs, such as FTP and TELNET, TCP and IP can be considered separate *layers*, each of which calls on the services of the layer below it. Generally, TCP/IP uses four layers:

- The application layer

This is the layer containing user processes cooperating with another process on the same or a different host. Examples are TELNET (a protocol for remote terminal connections), FTP (file transfer protocol) and SNMP (simple network management protocol). For details, see 2.4.5, “The Socket Application Programming Interface” on page 56.

- The transport layer

This is a layer such as TCP that provides reliable, *connection-oriented* transport services needed by many applications. For details, see 2.4.1, “Transmission Control Protocol (TCP) Layer Details” on page 50.

As an alternative for applications that do not require all the functions offered by TCP, the simpler but unreliable and *connectionless*, transport functions of UDP can be used. For details, see 2.5.1, “User Datagram Protocol (UDP) Layer Details” on page 64.

- The internetworking layer

The internetworking layer is also called *network* layer. The most important protocol is the internet protocol (IP). IP provides the basic service of getting datagrams to their destination. It is a *connectionless* protocol that does not assume reliability from the lower layers. IP does *not* provide reliability, flow control or error recovery. These functions must be provided at a higher level, either at the transport layer by using TCP as the transport protocol, or at the application layer if UDP is used as the transport protocol. A message unit in an IP network is called an *IP datagram*. For details, see 2.4.2, “Internet Protocol (IP) Layer Details” on page 52.

- The network interface layer

The network interface layer comprises the *data link control (DLC)* and the *physical* layer. This interface may or may not provide reliable delivery and may be packet or stream-oriented. In fact, TCP/IP does not specify any protocol here, but can use almost any network interface available, which illustrates the flexibility of the IP layer.

The data link layer contains the protocols needed to manage a specific physical medium, such as token-ring or a point-to-point line. For details, see 2.4.3, “Data Link Control Layer Details” on page 54. The physical layer defines the mechanical and electrical characteristic allowing a device to attach to a network. For details, see 2.4.4, “Physical Layer Details” on page 55.

TCP/IP networks are generally a large number of independent networks connected together by routers. Stations should be able to access computers or other resources on any of these networks. Datagrams will often pass through multiple different networks before getting to their final destination. The routing needed to accomplish this is invisible to the user. As far as the user is concerned, all he needs to know in order to access another system is an *Internet address*.




Figure 22. Dotted Decimal IP Addresses

Internet protocol (IP) addresses are 32-bit numbers. IP addresses are normally written as four decimal numbers, each representing eight bits of the address, separated by dots. Therefore we use the term *dotted decimal notation*. Routing decisions within an Internet are always based on the IP address of the destination node. It is important to realize that an IP address does not identify a station, but rather a specific network interface. For details, see 2.7, “IP Addresses” on page 71.

Generally the structure of the address gives you some information about how to get to the system. For example, 9 is a network number assigned by the central Internet authority to IBM. IBM uses the next byte to indicate for which country the address range is used. For example, 9.132 has been assigned to the Netherlands. The third byte indicates a physical network in the Netherlands. For example, 9.132.56 happens to be a token-ring used by the Zoetermeer IBM Global Network (IGN) technical support organization. The last byte allows for up to 254 systems on

each token-ring. (It is 254 because 0 and 255 are not allowed for reasons that will be discussed later.)

People tend to remember names more easily than numbers. Therefore, TCP/IP has implemented a procedure that allows users to refer to systems by name, rather than by Internet address. When we specify a name, the network software (referred to as the *name server*) looks it up in a database and comes up with the corresponding Internet address. It is essential that all routing within an Internet is done using the 32-bit destination address, and the name server technology is just an easy way to refer to systems.

The IP layer is built on a *connectionless* technology. Information is transferred as a sequence of *datagrams*. A datagram is a collection of data that is sent as a single message. Each of these datagrams is sent through the network individually. Individual packets may be routed along different routes, and datagrams may be received in a different order as sent. TCP, which is running on top of IP, provides methods to establish a connection, or conversation, that will continue for some time. TCP provides resequencing for datagrams that have been received out of order.




Figure 23. Fragmentation and Packetization

Figure 23 depicts the transfer of a 3000-byte file. For networks (for example, Ethernet) that cannot handle a 3000-byte datagram, the IP protocol will break this up into multiple smaller datagrams. Each of these datagrams will be sent to the other end. At that point, they will be put back together into the 3000-byte file. However, while those datagrams are in transit, the network does not know that there is any connection between them. It is perfectly possible that the second datagram will actually arrive before the first one. It is also possible that somewhere in the network, an error will occur, and some datagram will not get through at all. In that case, that specific datagram has to be sent again.

Note by the way that the terms *datagram* and *packet* often seem to be nearly interchangeable. Technically, datagram is the right word to use when describing TCP/IP. A datagram is a unit of data, which is what the protocols deal with. A packet is a physical thing, appearing on a token-ring or some wire. In most cases a packet simply contains a datagram, so there is very little difference.

However, they can differ. When TCP/IP is used on top of X.25, the X.25 interface may break the datagrams up into smaller packets. This process is referred to as *packetization*. This is invisible to IP, because the packets are put back together into a single datagram at the other end before being processed by TCP/IP. So in this case, one IP datagram would be carried by several packets. However, with

most media, there are efficiency advantages to sending one datagram per packet, and so the distinction tends to vanish.

Packetization does not take place on any of the network interfaces (frame relay, PPP, ESCON, and token-ring) currently supported by the 3746 IP router, and the terms datagram and packet can be used interchangeably. Note that on frame relay, PPP, and token-ring networks the term *frame* is more common than the term packet.

2.4.1 Transmission Control Protocol (TCP) Layer Details

Two separate protocols are involved in handling TCP/IP datagrams. TCP (the *transmission control protocol*) is responsible for breaking up (fragmenting) the message into datagrams, reassembling them at the other end, resending anything that gets lost, and putting things back in the right order. IP (the *internet protocol*) is responsible for routing individual datagrams. The IP layer on intermediate (IP forwarding) nodes may decide to fragment datagrams as well. However, reassembly is only done at end nodes.




Figure 24. Transmission Control Protocol. Processes X and Y communicate over a TCP connection carried by IP datagrams.

It may seem like TCP is doing all the work. However, in a larger Internet, getting a datagram to its destination can be a complex job. Keeping track of the routes to all of the destinations and handling incompatibilities among different transport media turns out to be a complex job. These tasks are the responsibility of the IP layer.

The interface between TCP and IP is fairly simple. TCP simply hands IP a datagram with a destination. IP delivers, using the services of the data link control layer, the datagram directly to its destination or forwards it to the next router in the direction of the destination. IP does not know how a datagram relates to any datagram before it or after it.

TCP/IP stations are identified by their IP addresses. The data exchanged between two stations employed in a TCP connection is routed through the Internet based on the IP address of the receiving station (the destination address).

Usually a single address is allocated per station, meaning that to keep track of multiple connections between two stations, TCP has to know which connection a datagram is part of. This task is referred to as *demultiplexing*. In fact, there are several levels of demultiplexing going on in TCP/IP.




Figure 25. TCP Connection Multiplexing

The information needed to do this demultiplexing is contained in a TCP *header*, which is added to the beginning of a TCP message. This header contains at least 20 bytes.




Figure 26. TCP Message

The source and destination *port numbers* are used to keep track of different conversations. Suppose three different people are transferring files. Your TCP might allocate port numbers 1000, 1001, and 1002 to these transfers. When you are sending a datagram, this becomes the *source* port number, since you are the source of the datagram. Of course the TCP at the other end has assigned a port number of its own for the conversation. Your TCP has to know the port number used by the other end as well. Which ports TCP uses will be explained later.

It puts this in the *destination* port field. Of course if the other end sends a datagram back to you, the source and destination port numbers will be reversed, since then it will be the source and you will be the destination.

Each datagram has a *sequence number*. This is used so that the other end can make sure that it gets the datagrams in the right order, and that it has not missed any. TCP does not number the datagrams, but the bytes. So if there are 500

octets of data in each datagram, the first datagram might be numbered 0, the second 500, the next 1000, the next 1500, etc.

To detect transmission errors a *checksum* is included in the TCP header. The checksum is a number that is based on the contents of the datagram. TCP at the other end computes the checksum again. If the receiving TCP disagrees with the sending TCP, transmission errors have occurred and the datagram is thrown away.

The remaining items in the header are generally involved with managing the connection. In order to make sure the datagram has arrived at its destination, the recipient has to send back an *acknowledgement*. This is a datagram whose *acknowledgement number* field is filled in. For example, sending a packet with an acknowledgement of 1500 indicates that you have received all the data up to byte number 1500. If the sender does not get an acknowledgement within a reasonable amount of time, it sends the data again.

A window is used to control how much data can be in transit at any one time. It is not practical to wait for each datagram to be acknowledged before sending the next one. That would slow things down too much. On the other hand, it would be unacceptable if stations just keep sending data, as this might overrun the capacity of the receiving stations. Furthermore, if one or multiple routers are involved in the transmission, uncontrolled sending of data might lead to the overrunning of the intermediate routers, leading to congestion and performance degradation.

Each end indicates how much new data it is currently prepared to absorb by putting the number of bytes in its *window* field. As the computer receives data, the amount of space left in its window decreases. When it goes to zero, the sender has to stop. As the receiver processes the data, it increases its window, indicating that it is ready to accept more data. Often the same datagram can be used to acknowledge receipt of a set of data and to give permission for additional new data (by an updated window).

The *urgent* field allows one end to tell the other to skip ahead in its processing to a particular byte. This is often useful for handling asynchronous events (for example, when you type a control character or other command that interrupts output). The other fields are beyond the scope of this document.

2.4.1.1 3746 IP and TCP

The 3746 IP router uses TCP for its TELNET, SNMP relay, and BGP functions.

2.4.2 Internet Protocol (IP) Layer Details

TCP is dependent on the functions of IP. When the TCP layer hands a datagram over to IP, it tells the IP layer the Internet address of the computer at the other end. This is all IP is concerned about. IP does not care about what is in the datagram, or even in the TCP header. IP's job is simply to find a route for the datagram and get it to the other end.

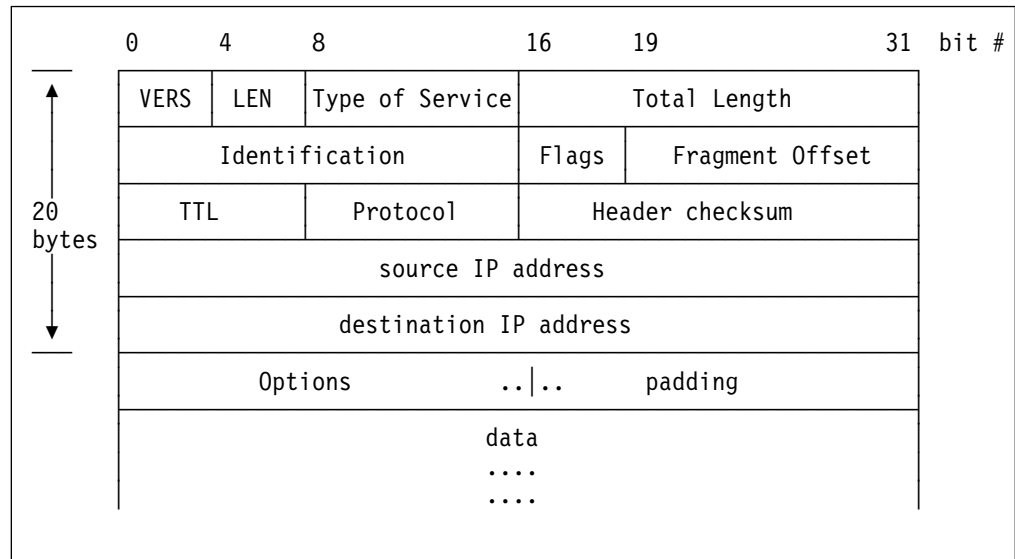


Figure 27. IP Datagram Format

In order to allow routers or other intermediate systems to forward the datagram, it adds its own IP header. The main things in this header are the source and destination Internet address (32-bit addresses, such as 9.132.56.80), the protocol number, and another checksum. The source Internet address is simply the address of your machine. It is required so the other end knows where the datagram came from. The destination Internet address is the address of the other machine. It is necessary so any routers in the middle can route to its final destination.

The *protocol number* tells IP at the other end to send the datagram to TCP. Although most IP traffic uses TCP, there are other protocols that can use IP, so you have to tell IP which protocol to send the datagram to. Finally, the *checksum* allows IP at the other end to verify that the header was not damaged in transit.

Note: TCP and IP have separate checksums. IP needs to be able to verify that the header did not get damaged in transit, or it could send a message to the wrong place. It is considered more efficient and safer to have TCP compute a separate checksum for the TCP header and data.

The header contains some additional fields that have not been discussed. The flags and fragment offset are used to keep track of the pieces when a datagram has to be split up. This can happen when datagrams are forwarded through a network for which they are too big. In addition to the fragmentation that is done on the station from which the data originates, IP routers along the route may decide to fragment the datagram as well (for example, when the size of the datagram exceeds the maximum transmission size of the network on which the datagram has to be forwarded).

Note: Reassembly of fragmented datagrams is only done at the receiving end node. No reassembly will be done on intermediate nodes.

The *time to live (TTL)* is a number that is decremented whenever the datagram passes through a system. When it goes to zero, the datagram is discarded. The TTL prevents IP packets from endless looping in a network. The Traceroute application, which enables a user to find the route IP packets take through the network, is based on the concept that IP packets with a TTL value of 1 are discarded (and an ICMP error message returned).

Note: Of course, in well-designed networks loops do not exist, but a proper network architecture has to cope with *impossible* conditions.

At this point, it is possible that no more headers are needed. If your computer happens to have a direct phone line connecting it to the destination computer, or to a router, it may simply send the datagrams out on the line (though likely a synchronous protocol such as HDLC would be used, and it would add at least a few bytes at the beginning and end).

2.4.2.1 3746 IP and IP Layer

As can be seen in Figure 27 on page 53, the IP header provides an option field. The 3746 IP router supports the following options in received packets:

- End of option list
- No operation
- Loose source route and strict source route

The processing for these options is similar; the primary difference is that the next destination IP address specified in the strict source route must be on one of the directly connected networks, while the next destination IP address specified in the loose source route may be on a remote network.

The 3746 does not originate IP datagrams with a source route option.

- Record route

The 3746 IP router updates the record route option field (if present) in forwarded packets only; it does not update this option for locally delivered datagrams.

The 3746 does not originate IP datagrams with a record route option. 3746 ICMP echo request (PING) messages do not contain this option.

- Internet time stamp

The 3746 IP router supports three forms of time stamp options:

- Time stamps only
- IP address and time stamps
- IP address and time stamps with the IP addresses prespecified

The 3746 IP router updates the time stamp option field (if present) in forwarded packets only, it does not update this option for locally delivered datagrams.

The 3746 IP router recognizes but ignores the security and stream ID options. It stops processing when encountering an option it does not recognize.

2.4.3 Data Link Control Layer Details

The protocols that define information exchange over a network are known as data link layer protocols.

It is mandatory that they be adhered to by all devices, whether end or intermediate nodes, that attach to a network. Data link layer definition of protocols define the procedures for devices to transfer data over a network. They cover aspects such as:

- The format of data link layer packets
- The format of data link layer addresses

- The process for transmitting and receiving packets
- Connection establishment and termination requirements
- Error detection and handling

2.4.3.1 3746 IP and DLC Layer

Examples of data link control protocols supported by the 3746 IP router are:

- The point-to-point protocol (PPP)

The 3746 IP router provides PPP using leased lines over V.24, V.35 and X.21. The PPP protocol is described in RFC 1331 and 1332.

- Frame relay

The 3746 IP router provides frame relay support using leased lines over V.24, V.35 and X.21. IP packets are sent using the encapsulation methods described in RFC 1490. The ITU and ANSI frame relay standards are detailed in Q.922/Q.923 and T1.617/T1.618, respectively.

- Token-ring

The 3746 provides IP attachments using IEEE 802.5 (token-ring). The DLC used on token-ring is specified in IEEE 802.2.

- ESCON Channel Data Link Control (CDLC)

The 3746 IP router provides IP over CDLC to connect to TCP/IP V3R1 for MVS, OS/390 V2R4 TCP/IP stack (not supported by the open edition stack), and OS/390 V2R5 TCP/IP stack.

The IEEE 802 LAN protocols differ from the others above in that the data link layer functions are separated into two distinct sublayers. The medium access control (MAC) sublayer defines how devices should access and send packets over a LAN. The data link sublayer defines how information exchange should take place between devices on an individual LAN.

2.4.4 Physical Layer Details

Although TCP/IP does not identify the physical layer as a separate layer, the DLC layer assumes its presence. The protocols that define physical device attachment are known as the physical layer protocols. It is mandatory that they be adhered to by all devices that attach to a physical network.

Physical layer protocols define the mechanical and electrical characteristics for a device to attach to a network. They cover aspects such as:

- Attachment mechanisms (for example, connector types and pinouts)
- Electrical characteristics (for example, interface voltages and impedances)

Physical and data link layer protocols are specified such that physical layer implementations can support a range of alternate data link protocols, and data link protocols can operate over a range of physical layer implementations. This advantage results from the layered structure of the OSI reference model. An example of this would be ITU X.25 data link layer protocols, which can operate over V.24 and V.35 physical layers.

2.4.4.1 3746 IP and Physical Layer

Examples of physical layer protocols supported by the 3746 IP router are:

- ESCON
- ITU X.21, V.24, and V.35
- IEEE 802.5 (token-ring)

Note: Ethernet support is provided by bridging a 3746 token-ring port to an Ethernet LAN (see 5.6, “Ethernet” on page 175).

2.4.5 The Socket Application Programming Interface

So far, we have described how a stream of data is broken up into datagrams, sent to another computer, and put back together. However, something more is needed in order to accomplish anything useful. There has to be a way for you to open a connection to a specified computer, log into it, tell it what file you want, and control, for example, the transmission of the file. This is done by *application protocols*.

The application protocols run *on top* of TCP/IP. That is, when they want to send a message, they give the message to TCP. TCP makes sure it gets delivered to the other end. Because TCP and IP take care of all the networking details, the application protocols can treat a network connection as if it were a simple byte stream.

To enable the exchange of data between application programs, a method is required for how a connection between applications can be established. Suppose you want to send a file to a computer whose Internet address is 9.132.56.80. To start the process, you need more than just the Internet address. You have to connect to the FTP server at the other end. In general, network programs are specialized for a specific set of tasks. Most systems have separate programs to handle file transfers, remote terminal logins, mail, etc. When you connect to 9.132.56.80, you have to specify that you want to talk to the FTP server. This is done by having *well-known sockets* for each server.

TCP uses port numbers to keep track of individual conversations. User programs normally use more or less random port numbers. However, specific port numbers are assigned to the programs that sit waiting for requests. For example, if you want to send a file, you will start a program called *FTP*. It will open a connection using some random number, say 1234, for the port number on its end. However, it will specify port number 21 for the other end. This is the official port number for the FTP server.

Two different programs are involved. The FTP client is the program designed to accept commands from your terminal and pass them on to the other end. The program that you talk to on the other machine is the FTP server. It is designed to accept commands from the network connection, rather than an interactive terminal.

There is no need for the FTP client to use a well-known socket number for itself. Nobody is trying to find it. However, the servers have to have well-known numbers, so that people can open connections to them and start sending them commands. Several port numbers have been reserved for specific programs. Reserved numbers are listed in RFC 1700 - *Assigned Internet Numbers* (see 2.12, “Request for Comments (RFCs)” on page 90).




Figure 28. Sockets and Ports

Note that a connection is actually described by a set of four numbers: the Internet address at each end and the TCP port number at each end. Every TCP datagram has all four of those numbers in it. (The Internet addresses are in the IP header, and the TCP port numbers are in the TCP header.) In order to keep things straight, no two connections can have the same set of numbers. However, it is enough for any one number to be different. Figure 28 depicts a situation where two different users are sending files to the same other machine.

Since the same machines are involved, the Internet addresses are the same. Since they are both doing file transfers, one end of the connection involves the well-known port number for FTP. The only thing that differs is the port number for the program that the users are running. That's enough of a difference. Generally, at least one end of the connection asks the network software to assign it a port number that is guaranteed to be unique. Normally, it is the user's end, since the server has to use a well-known number.

2.4.6 Application Defaults

An extensive description of the TCP/IP application protocols is outside the scope of this document. The main function of the 3746 IP router is IP forwarding. The only applications relevant to the 3746 IP router are the SNMP agents functions, the (TELNET) functions for remote configuration and operation, and the dynamic routing protocols (RIP, OSPF, and BGP).

An understanding of the TCP/IP port number concept is sufficient to configure the 3746 IP, the access control and filtering functions. For readers that require more information, we advise you to read the RFCs (see 2.12, "Request for Comments (RFCs)" on page 90) describing the applications.

Some general remarks can be made. TCP/IP is intended to be usable on any computer. Unfortunately, not all computers agree on how data is represented. There are differences in character codes (ASCII vs. EBCDIC), in end-of-line conventions (carriage return, line feed, or a representation using counts), and in whether terminals expect characters to be sent individually or a line at a time. In order to allow computers of different kinds to communicate, each application's protocol defines a standard representation. TCP and IP do not care about the representation. TCP simply sends bytes. However, the programs at both ends have to agree on how the bytes are to be interpreted. The RFC for each application specifies the standard representation for that application.

2.4.7 IP Routing

It has been mentioned before that the IP layer is responsible for getting datagrams to the station indicated by the destination address. The task of finding how to get a datagram to its destination is referred to as *routing*. In fact, many of the routing details depend upon the particular implementation. However, some general things can be said.

First, it is necessary to understand the model on which IP is based. Each station has to be attached to at least one network. Most often this will be a local area network attachment, but stations might be attached via a serial line (PPP or frame relay) or, for that matter, be ESCON attached.




Figure 29. Types of Networks

Stations can send datagrams to any other system in its own network. Networks can be:

- Multi-access broadcast networks

Multi-access broadcast networks allow the attachment of multiple stations. To send data to a destination node attached to the same network, the physical address of the node must be learned, and direct communication can take place. Broadcasting allows a node to send a datagram to multiple stations at the same time.

Examples of multi-access broadcast networks are token-ring, Ethernet, and FFDI.

- Multi-access non-broadcast networks

Multi-access non-broadcast networks also allow the attachment of multiple stations. To send data to a destination node attached to the same network, the physical address of the node must be learned, and direct communication can take place. In general, unless special provisions are taken, broadcasting is not possible.

Examples of multi-access non-broadcast networks are frame relay and X.25.

- Point-to-point (PtP) networks

Point-to-point networks allow direct communication between two adjacent nodes. No broadcast is required on PtP connections. From an IP routing perspective each PtP connection is a separate network, although only two

stations connect. Instead of the term *PtP network*, the term *PtP connection* is more commonly used.

Examples of point-to-point networks are PPP, SLIP, SNAlink, and ESCON connections.

Stations that are attached to the same network can communicate directly. A problem arises when a system wants to send a datagram to a system on a different network. This problem is handled by routers. A router is a system that connects a network with one or more other networks. Routers are often normal computers that happen to have more than one network interface.




Figure 30. IP Routing

Figure 30 depicts an example of a 3746 IP controller that has two token-ring interfaces and an ESCON connection. Thus it is connected to networks 9.132.4, 9.132.5, and 9.132.6. The 3746 IP can act as a router between the networks, allowing communication between the stations attached to any of the networks.

Routing in IP is based entirely upon the network number of the destination address. Each computer has a table of network numbers. For each network number, a router is listed. This is the router to be used to get to that network. Note that the router does not have to connect directly to the network. In the previous example, (frame relay) network 9.132.7 is not directly attached, and datagrams originating from the MVS host destined for a frame relay-attached host must be forwarded to another IP router first.

When a computer wants to send a datagram, it first checks to see if the destination address is on the system's own local network. If so, the datagram can be sent directly. Otherwise, the system expects to find an entry for the network that the destination address is on. The datagram is sent to the router listed in that entry.

Various strategies have been developed to reduce the size of the routing table. One strategy is to depend upon *default routes*. Often, there is only one router out of a network. This router might connect a local token-ring to a campus-wide backbone network. In that case, we do not need to have a separate entry for every network in the world. We simply define that router as the *default*. When no specific route is found for a datagram, the datagram is sent to the default router. A default gateway can even be used when there are several routers on a network.

There are provisions for routers to send a message saying *I'm not the best router -- use this one instead.*, (the message is sent via ICMP).

Note: See the description of the ICMP redirect message in 2.5.2.1, "3746 IP and ICMP" on page 65.

Most network software is designed to use these messages to add entries to their routing tables. As can be seen, 9.132.56 has two routers, 9.132.56.2 and 9.132.56.3. 9.132.56.2 leads to several other IBM-internal networks, while 9.132.56.3 leads indirectly to all external (non-IBM) networks.




Figure 31. ICMP Redirect

Suppose host A attached to network 9.132.56 does not have any routing table entries other than router 9.132.56.2 that has been defined as a default router. Now what happens when A needs to send a datagram to host B. Since A has no entry for external networks, the datagram will be sent to the default router, 9.132.56.2. As it happens, this router is the wrong one. So it will forward the datagram to 9.132.56.3. But it will also send back an error saying, in effect, *to get to the external network, use 9.132.56.3*. The host station will then add an entry to the routing table and any future datagrams will then go directly to 9.132.56.3. (The error message is sent using the ICMP protocol. The message type is called *ICMP redirect*.) The 3746 IP router supports sending ICMP redirect messages, but ignores any ICMP redirect messages that it receives.

It is recommended that individual computers do not try to keep track of the entire network. Instead, they should start with default routers and let the routers tell them the routes, as just described. However, this does not say how the routers should find out about the routes.

Routers cannot depend upon this strategy. They have to have fairly complete routing tables. For this, a dynamic routing protocol is needed. A routing protocol is simply a technique for the routers to find each other and keep up to date about the best way to get to every network. An excellent source of information on routing protocols is *Routing in the Internet*, SR28-5714. The routing protocols supported on the 3746 IP router are RIP, OSPF, and BGP. They are described in detail in Chapter 3, "Dynamic IP Routing Protocols - Introduction" on page 91.

2.4.7.1 Direct and Indirect Destinations

If the destination host is attached to a network to which the source host is also attached, an IP datagram can be sent directly, simply by encapsulating the IP datagram in the physical network frame. This is called *direct routing*.

Indirect routing occurs when the destination host is not on a network directly attached to the source host. The only way to reach the destination is via one or more routers. The address of the first of these routers (the *first hop*) is called an *indirect route*. The first hop address is the only information needed by the source host; the router which receives a datagram has responsibility for the second hop and so on.




Figure 32. Direct and Indirect IP Routes. Host A has a direct route to hosts B and D, and an indirect route to host C. Host D is a router between the 129.1 and 129.2 networks.

A host can tell whether a route is direct or indirect by examining the network number and subnet number parts of the IP address. If they match one of the IP addresses of the source host, the route is a direct one.

The host needs to be able to address the target correctly using a lower-level protocol than IP. This can either be done automatically using a network-specific protocol, such as ARP (see 2.9, “Address Resolution Program (ARP)” on page 81), which is used on broadcast LANs, or by statically configuring the host.

2.4.7.2 IP Routing Table

Each host keeps the set of mappings between destination IP addresses and the IP addresses of the next hop routers for those destinations in a table called the *IP routing table*.

Three types of mappings can be found in this table:

1. Direct routes, for locally attached networks
2. Indirect routes, for networks reachable via one or more routers
3. A default route, which contains the IP address of a router to be used for all IP addresses that are not covered by the direct and indirect routes

See the network in Figure 33 on page 62 for an example configuration.




Figure 33. Example IP Routing Table

The routing table of host D will contain the following entries:

Destination Route via

128.10 direct attachment

128.15 direct attachment

129.7 128.15.1.2

default 128.10.1.1

2.4.7.3 IP Routing Algorithm

From the foregoing discussion, one can easily derive the steps that IP must take in order to determine the route for an outgoing datagram. This is called the *IP routing algorithm* and it is shown schematically in Figure 34.

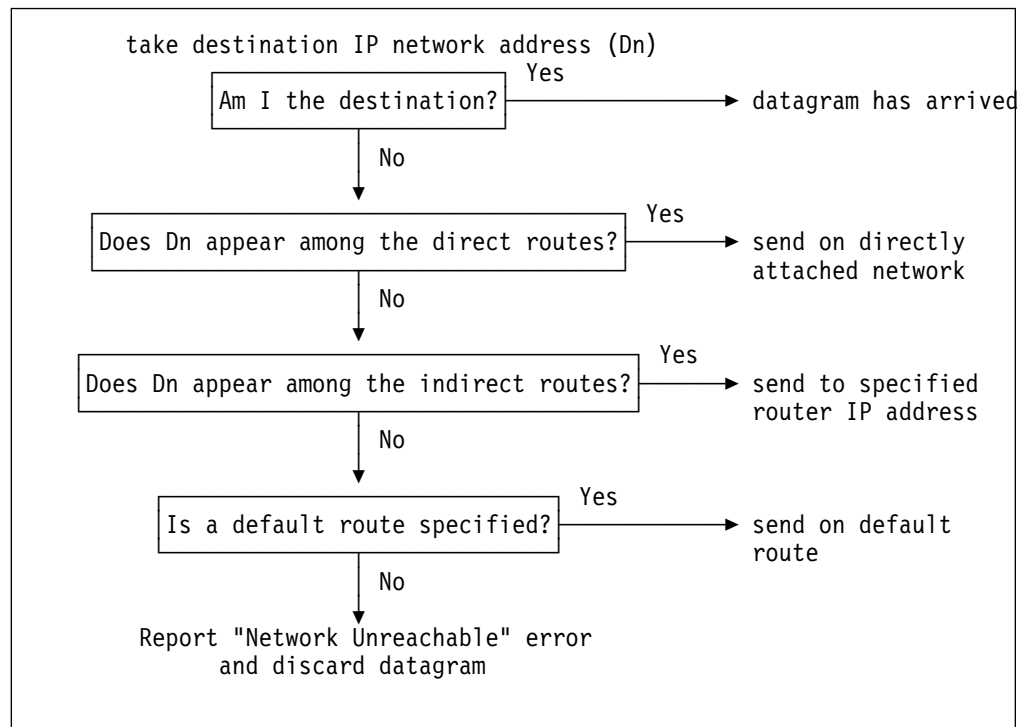


Figure 34. IP Routing Algorithm

Note that this is an iterative process. It is applied by every host handling a datagram, except for the host to which the datagram is finally delivered.

2.4.7.4 Multicasting

The 3746 IP router supports local delivery and forwarding of datagram addresses to multicast (class D) IP addresses. It provides an interface by which multicast addresses can be registered for local delivery. The 3746 IP router supports both local delivering and forwarding a multicast datagram. It never forwards datagrams addressed to a local multicast IP address (224.0.0.x). Local multicast IP address 224.0.0.4 and 224.0.0.5 are used by OSPF.

2.5 Protocols Other Than TCP

So far, we have described only connections that use TCP. Recall that TCP is responsible for breaking up messages into datagrams and reassembling them properly. However, in many applications, we have messages that will always fit in a single datagram. An example is SNMP. SNMP queries are in general very short and will fit in one datagram. So will the answer. Thus it is overkill to use TCP.




Figure 35. UDP

Of course, TCP does more than just break things up into datagrams. It also makes sure that the data arrives, resending datagrams where necessary. But for a question that fits in a single datagram and is not dependent on the reliability offered by TCP, the complexity of TCP is not needed. If we do not get an answer after a few seconds, we can just ask again. For applications like this, there are alternatives to TCP.

2.5.1 User Datagram Protocol (UDP) Layer Details

The most common alternative to TCP is UDP (*user datagram protocol*). UDP is designed for applications where you do not need to put sequences of datagrams together. It fits into the system much like TCP. There is a UDP header. The network software puts the UDP header on the front of your data, just as it would put a TCP header on the front of your data. Then UDP sends the data to IP, which adds the IP header, and puts UDP's protocol number (17) in the protocol field instead of TCP's protocol number.




Figure 36. UDP Datagram Format

However, UDP does not do as much as TCP does. It does not split data into multiple datagrams. It does not keep track of what it has sent so it can resend if necessary. About all that UDP provides is port numbers, so that several programs can use UDP at once. UDP port numbers are used just like TCP port numbers. There are well-known port numbers for servers that use UDP. Note that the UDP header is shorter than a TCP header. It still has source and destination port numbers and a checksum, but that is about it. There is no sequence number, since it is not needed. Examples of applications that use UDP are SNMP and name server lookups.

2.5.1.1 3746 IP and UDP

The 3746 IP router uses UDP for RIP and SNMP data transport.

If a UDP packet with an unknown (unregistered) destination port is received, this event is logged. For destination ports other than 9 (RFC 863, *Discard protocol*) an ICMP destination unreachable message (port unreachable) is returned, unless the packet was a link-level broadcast or the source IP address was a broadcast or multicast IP address.

2.5.2 Internet Control Message Protocol (ICMP)

Another alternative protocol is ICMP (*Internet control message protocol*). ICMP is used for error messages and other messages intended for the TCP/IP software itself, rather than any particular user program. For example, if you attempt to connect to a host, your system may get back an ICMP message saying *host unreachable*.

ICMP can also be used to find out some information about the network. ICMP is similar to UDP in that it handles messages that fit in one datagram. However, it is even simpler than UDP. It does not even have port numbers in its header. Since all ICMP messages are interpreted by the network software itself, no port numbers are needed to say where an ICMP message is supposed to go.

ICMP messages are sent in IP datagrams. The IP header will always have a protocol number of 1, indicating ICMP and a type of service of zero (routine). The IP data field will contain the actual ICMP message in the format shown in Figure 37.




Figure 37. ICMP Message Format

2.5.2.1 3746 IP and ICMP

The 3746 IP router recognizes the following ICMP types (see the Type field in Figure 37):

Type Description

0 Echo Reply

When an Echo request is received, an Echo reply is returned containing the data from the Echo request.

3 Destination Unreachable

The 3746 may generate the following ICMP destination unreachable messages:

- Network unreachable

When there is no active route to the destination address, the destination network is zero, or the destination address is not a valid class A, B, C, or D address.

- Fragmentation need and DF set

An outgoing packet needs fragmentation, but the DF bit in its IP header is set, or fragmentation results into an invalid fragment offset.

- Protocol unreachable

An IP datagram destined for the 3746 IP router contains an unsupported protocol number.

- Port unreachable

A UDP datagram destined for the 3746 IP router contains an unsupported UDP port number.

- Source route failed

The next hop IP address in a datagram strict source routing is not on a directly connected network.

4 Source quench

Source quench messages are logged but no action is taken. The 3746 never generates a source quench message.

5 Redirect

Redirect messages are logged but no action is taken. The 3746 generates a redirect message (redirect datagram for the host) when a packet is forwarded on the same interface as received, the network is not point-to-point, and the subnet or network number of the source IP address matches that of the next hop IP address.

8 Echo Request

Echo request messages are responded to with the echo reply message. ICMP Echo requests are generated by the PING application (see "PING" on page 67) at the user's request.

11 Time exceeded

Time exceeded messages are logged. The 3746 generates a time exceeded message (time to live exceeded in transit) if the time to live (TTL) field in a datagram the 3746 IP router is forwarding is 0 or 1.

12 Parameter problem

Parameter problem messages are logged. The 3746 generates a parameter problem message (pointer indicates the error) if:

- Options field in IP header indicates that the options extend beyond the end of the IP header.
- Pointer is less than 4 in loose source route, strict source route, or record route option. Pointer is less than 5 in a time stamp option.
- Pointer is less than the length in a record route or time stamp option, but insufficient room for the IP router to insert its IP address and/or time stamp.
- In a record route option, the option is full (no data can be added).

13 Time stamp

Time stamp messages are logged. The 3746 never generates or responds to time stamp messages.

14 Time stamp reply

Time stamp reply messages are logged. The 3746 never generates or responds to time stamp reply messages.

15 Information request

Information request messages are logged. The 3746 never generates or responds to information request messages.

16 Information reply

Information reply messages are logged. The 3746 never generates information reply messages.

17 Address mask request

Address mask request messages are dropped if:

- The destination IP address matches the 3746's internal IP address
- The source IP address is 0.0.0.0 and the receiving interface has more than one IP address assigned to it
- The source IP address is 0.0.0.0 and the receiving network is not broadcast or point-to-point

The 3746 never generates address mask request messages.

18 Address mask reply

Address mask reply messages are returned in response to valid address mask request messages. An address mask reply message is returned with the destination IP address set to 255.255.255.255 if the (request) source IP address is 0.0.0.0; otherwise, the (reply) destination IP address is set to the (request) source IP address.

When the 3746 IP router receives an ICMP message with a Type value it does not recognize, it logs the event. If the Type is greater than 18, it increments the count of erroneous received ICMP messages. The 3746 keeps transmit and receive counts per-Type for values less than 19.

The TTL value in the IP header of the ICMP messages generated by the 3746 IP router is equal to the value that is generally set for all IP datagrams, namely 60. Precedence and TOS fields are set to zero. No options are included. No mechanism exists to limit the rate at which ICMP messages are sent.

2.5.2.2 PING and Traceroute

Two important applications implemented on the 3746 IP router that rely on ICMP messages are PING and Traceroute. Both applications are extremely useful in diagnosing connectivity problems. Each application can be invoked by using either CCM or in TELNET command line mode.

PING: PING is the simplest of all TCP/IP applications. It sends one or more messages to a specified destination host requesting a reply and measures the round trip time. The word *PING*, which is used as a noun and a verb, is taken from the sonar operation to locate an underwater object. It is also an abbreviation for *packet internet groper*.

Traditionally, if you could PING a host other applications, such as TELNET or FTP could reach that host. With the advent of security measures on the Internet, which control access to networks by application protocol and/or port number, this is no

longer strictly true. Nonetheless, the first test of reachability for a host is still to attempt to PING it.




Figure 38. Packet InterNet Groper (PING)

The syntax that is used in different implementations of PING varies from platform to platform. The syntax here is for the 3746 IP implementation:

PING host

where host is the target IP address. On the 3746 IP router, the target IP address specified must be in dotted decimal notation. Host names are not supported.

PING uses the ICMP Echo and Echo Reply messages. Since ICMP is required in every TCP/IP implementation, hosts do not require a separate server to respond to PINGs.

PING is useful for verifying IP connectivity (see also 6.3.1, “Using PING” on page 267). Be aware that IP is a connectionless protocol and that a PING request and its reply may traverse different routes through the network. Being able to send a PING to a node does not necessarily mean that the reply can be returned as well. PING requests contain a destination address that is set to the target host indicated in the PING command. The 3746 IP routers set the source IP address to its *router_ID* (see 2.7.4, “3746 IP and IP Addresses” on page 79). Therefore, to enable the target host to return the PING reply, an active route between this host and the 3746 Router_ID is required.

Note: For PINGs to be successful, make sure that the PING ICMP messages are not filtered along the route between the 3746 IP router and the destination host.

In normal operation PING builds and sends an ICMP echo request. PING waits up to 1 second to receive an ICMP echo reply to calculate the round-trip time, display the size of the echo request/reply packet (not including IP header), its source IP address, and the round-trip in milliseconds. 3746 IP echo request/reply messages are always 64 bytes long (excluding IP header).

PING keeps count of the echo requests it sends and the echo replies it receives. It keeps a sum of the round-trip times and calculates maximum and minimum round-trip times. In addition, it keeps count of the ICMP destination unreachable (network or host unreachable). 3746 IP PING stops when the user presses any key. PING statistics are then displayed which include:

- The number of echo requests sent
- The number of echo replies received
- The percentage of packets lost (no replies received)

- The number of destination unreachable messages received
- The number of echo requests not sent due to lack of buffers

Traceroute: The PING command described in the previous section can be used to verify connectivity between two hosts. In addition, when the PING command fails, the Traceroute command can be used to determine the route that IP datagrams follow from host to host.

Traceroute relies on ICMP time exceeded and destination unreachable messages. It sends a UDP datagram with a TTL of 1 to the destination host. The first router to see the datagram will decrement the TTL to 0, return an ICMP time exceeded message, and will discard the datagram. In this way, the first router in the path is identified. This process can be repeated with successively larger TTL values in order to identify the series of routers in the path to the destination host. Traceroute sends UDP datagrams to the destination host, which reference a port number that is outside the normally used range. This enables Traceroute to determine when the destination host has been reached, that is, when an ICMP destination (Port) unreachable message is received.

For each TTL value, the UDP datagram is sent three times, waiting up to three seconds between transmissions to receive a response. Traceroute continues this process until either a TTL value of 33 has been reached or an ICMP destination unreachable message is received.




Figure 39. Traceroute

When a route exists between the 3746 IP router and a destination node, each intermediate router will return an ICMP time exceeded message. However, a situation might exist where somewhere along the route the Traceroute frame cannot be forwarded.

Note: For Traceroute to be successful, make sure that the Traceroute ICMP messages are not filtered along the route between the 3746 IP router and the destination host.

For each TTL value, Traceroute displays a line with the following information:

- 'BF' each time the 3746 IP is unable to send the UDP frame due to lack of buffers
- '*' each time no response is received

- The source IP address of the response, if different from the previous response received
- The round-trip time in milliseconds
- '!' if the response is a destination unreachable (port unreachable)
- '!P' if the response is a destination unreachable (protocol unreachable)
- '!N' if the response is a destination unreachable (network unreachable)
- '!H' if the response is a destination unreachable (host unreachable)

2.6 The Domain Name System

As indicated earlier, the network software generally needs a 32-bit Internet address in order to open a connection or send a datagram. However, users prefer to deal with computer names rather than numbers. Therefore, name servers can be accessed that keep track of host names and the corresponding Internet addresses. (In fact, these servers are somewhat more general than that. This is just one kind of information stored in the domain system.) A set of interlocking servers are used, rather than a single central one because there are so many different institutions connected to the Internet that it would be impractical for them to notify a central authority whenever they installed or moved a computer. Thus naming authority is delegated to individual institutions. The name servers form a tree, corresponding to institutional structure. The names themselves follow a similar structure.




Figure 40. Domain Name Server

A typical example is the name WITDER.BENELUX.IBM.COM. This is a computer at the IBM Zoetermeer networking center. In order to find its Internet address, you might potentially have to consult four different servers. First, you would ask a local domain server where the COM server is. COM is a server that keeps track of commercial institutions. The local domain server would give you the names and Internet addresses of several servers for COM. (There are several servers at each level to allow for the possibility that one might be down.) You would then ask COM where the server for IBM is. Again, it would give you names and Internet addresses of several servers for IBM. Then you would ask IBM where the server for BENELUX is, and finally you would ask one of the BENELUX servers about WITDER. The final result would be the Internet address for WITDER.BENELUX.IBM.COM.

Note: In practice there are quicker ways to get to the appropriate name server. However, detailing this is outside the scope of this document.

Each of these levels is referred to as a *domain*. The entire name, WITDER.BENELUX.IBM.COM, is called a *domain name*. (So are the names of the higher-level domains, such as BENELUX.IBM.COM, IBM.COM, and COM.)

Fortunately, you do not really have to go through all of this most of the time. First of all, the root name servers also happen to be the name servers for the top-level domains such as COM. Thus a single query to a root server will get you to IBM. Second, software generally remembers answers that it got before. So once we look up a name at BENELUX.IBM.COM, our software remembers where to find servers for BENELUX.IBM.COM, IBM.COM, and COM. It also remembers the translation of WITDER.BENELUX.IBM.COM. Each of these pieces of information has a *time to live* associated with it. Typically this is a few days. After that the information expires and has to be looked up again. This allows institutions to change things.

The domain system is not limited to finding out Internet addresses. Each domain name is a node in a database. The node can have records that define a number of different properties. Examples are Internet address, computer type, and a list of services provided by a computer. A program can ask for a specific piece of information, or all information about a given name. It is possible for a node in the database to be marked as an *alias* (or nickname) for another node. It is also possible to use the domain system to store information about users, mailing lists, or other objects.

There is an Internet standard defining the operation of these databases, as well as the protocols used to make queries of them. Every network utility has to be able to make such queries, since this is now the official way to evaluate host names. Generally utilities will talk to a server on their own system. This server will take care of contacting the other servers for them. This keeps down the amount of code that has to be in each application program.

2.6.1 3746 IP and the Domain Name System

The main function of the 3746 IP controller is to route IP traffic. To perform this function the 3746 IP has multiple IP network interfaces to which a single or multiple IP addresses have been assigned. To access the SNMP agent running on the 3746, one of its network interfaces must be referred to. The TELNET server is accessed by referring to the NNP's service LAN IP address.

The 3746 IP routing functions are fully independent from the naming convention used in your Internet. However, assigning a symbolic name to the 32-bit IP interface addresses simplifies the access to the 3746 network management and operation functions. No name server address needs to be configured in the 3746 IP router. 3746 IP can operate in networking environments that have no name server installed.

2.7 IP Addresses

As indicated earlier, Internet addresses are 32-bit numbers, normally written as four bytes (in decimal). An example is 9.132.6.2. IP addresses are comprised of a network identifier and a host identifier. IP addresses must be unique within an Internet. Hosts that are connected to the same network must have the same network identifier.




Figure 41. Host and Network Identifier

An IP address indicates both the network and the host within the network. Although the total number of IP addresses that can be used is high (2^{32}), this number is bound and care has been taken to enable an efficient allocation of addresses trying to solve the problem:

- Routers make routing decision based on the network identifier. Therefore, the network identifier must be unique throughout the network.
- Many networks exist and the number of networks is growing every day. The addressing structure must allow for an adequate number of networks.
- Large and small networks exist. LANs might connect thousands of hosts within the same network, while on a point-to-point connection only two stations connect.

Therefore, an addressing structure must exist that on one hand allows a large number of hosts within a single network, but on the other hand avoids a waste of addresses. To illustrate the latter, assume an addressing structure that has a 16-bit network and a 16-bit host identifier. If this structure is used for a point-to-point network (connecting only two hosts), only two out of 2^{16} addresses are used. The remaining addresses cannot be used elsewhere in the network.

The first solution to this problem is to use different classes of addresses whereby the length of the network and the host portion depend on the address class. The five different types of IP address classes that exist are depicted in Figure 42 on page 73. Address classes A to E are identified, whereby class D and class E addresses have special purposes. All IP addresses are fixed in length, namely 32 bits.

Note: Technically speaking, network number and address class prefix are different entities. However, when referring to a network number, the address prefix is considered part of it.

Class A addresses have a network number that is 8-bits long. The first bit of class A addresses is set to B'0'. As a network number of all zeroes cannot be used, the maximum number of class A networks is 2^7-1 , ranging from 1 to 127.

Note: Network 127 is used for *loopback*. For example, a host referring to 127.0.0.1 is referring to itself.




Figure 42. Assigned Classes of Internet Addresses

The maximum number of stations that can attach to a class A network is (2^{24}), (some host addresses have special meaning and cannot be used). Using class A addresses, very large networks can be built. IBM, for example, is using network 9 addresses for its internal network.

Class B addresses have a network number that is 16 bits long. The first 2 bits of class B addresses are set to B'10'. The maximum number of class B networks is 2^{14} . The first byte of class B addresses ranges from 128 to 191. The maximum number of stations that can attach to a class B network is (2^{16}), (some host addresses have special meaning and cannot be used). Class B addresses are typically used for medium-sized to large networks.

Class C addresses have a network number that is 24 bits long. The first 3 bits of class C addresses are set to B'110'. The maximum number of class C networks is 2^{21} . The first byte of class C addresses ranges from 192 to 223. The maximum number of stations that can attach to a class C network is (2^8), (some host addresses have special meaning and cannot be used). Class C addresses are typically used for small-sized networks.

Within an Internet (a collection of connected networks) IP addresses must be unique. A central Internet organization (the InterNIC) is responsible for allocating network addresses. The host numbers are assigned by the authority that owns the network.

To build an Internet an official IP address range has to be requested. Class A and class B addresses are no longer readily available, and only class C network numbers can be obtained. If the number of hosts within your network cannot be accommodated in a single class C network, multiple network numbers must be requested. Class C network numbers are allocated in sequence to allow address aggregation (see 2.7.3, "Supernetting (Route Aggregation)" on page 79).

Addresses must be unique within an Internet. However, within isolated IP networks the same addresses can be used. Certain network addresses have been reserved for general use. Organizations may decide to use one of these network addresses. However, when networks that are using the same address merge, addresses have to be changed. This process can be very cumbersome. It is therefore

recommended that you apply for a registered network number before building a private network.

Addresses above 223 are reserved for special purposes. Class D addresses are reserved for multicasting. Multicasting is for example, used by OSPF. For details, see 3.5.3, “Multicast Extensions to OSPF (MOSPF)” on page 115. Class E addresses are reserved for future use.

2.7.1 Subnetting

Identifying address classes has enabled a more efficient use of IP addresses; however, that is not sufficient. Remember that for each network a unique network number has to be used. IBM uses for its internal Internet network number 9, and unless additional techniques are used, a single point-to-point connection would use the whole address range. This would lead to a tremendous waste (about 2^{24}) of addresses, because they cannot be used elsewhere.




Figure 43. Subnetting

Therefore the concept of *subnetting* has been introduced. Instead of relying on the network/host subdivision implied by the address class, a *subnet mask* can be defined that subdivides each IP address host number into a *subnet* number and a host identifier. The more common name for subnet is *subnet*. The combination of the subnet number and the host identifier is often termed the *local address* or the *local part*. Network and subnet numbers together comprise the network identifier.




Figure 44. Address Subnetting

IP routing in a subnetted network requires an IP router to forward data from one subnet to another. Routing decisions are made using the network identifier of the destination IP address. Subnetting is transparent to remote networks. Stations in remote networks still regard the network number as the network identifier, and the host number as the host identifier.

The division of the local part of the IP address into subnet number and host identifier can be chosen freely by the local administrator; any bits in the local part can be used to form the subnet accomplished. The division is done using a *subnet mask*, which is a 32-bit number. Zero bits in the subnet mask indicate bit positions ascribed to the host identifier, and ones indicate bit positions ascribed to the subnet number. The bit positions in the subnet mask belonging to the network number are

set to one but are not used. Subnet masks are usually written in dotted decimal form, similar to IP addresses.

The special treatment of *all bits zero* and *all bits one* applies to each of the three parts of a subnetted IP address just as it does to both parts of an IP address that has not been subnetted. See 2.7, “IP Addresses” on page 71. For example, a subnetted class B network, which has a 16-bit local part, could use one of the following schemes:

- The first byte is the subnet number, and the second byte is the host identifier. This gives us 254 (256 minus 2 with the values 0 and 255 being reserved) possible subnets, each having up to 254 hosts. The subnet mask is 255.255.255.0.
- The first 12 bits are used for the subnet number and the last four for the host identifier. This gives us 4094 possible subnets (4096 minus 2) but only 14 hosts per subnet (16 minus 2). The subnet mask is 255.255.255.240.

There are many other possibilities.

While the administrator is completely free to assign the subnet part of the local address in any legal fashion, the objective is to assign a *number* of bits to the subnet number and the remainder to the host identifier. It is normal to use a contiguous block of bits at the beginning of the local address part for the subnet number because this makes the addresses more readable (this is particularly true when the subnet occupies 8 or 16 bits). With this approach subnet masks such as 255.255.255.0 and 255.255.255.240 are *good* masks, but masks such as 255.255.252.252 and 255.255.255.15 are not.

Using subnetting requires special care:

- Hosts attached to the same subnet must use the same subnet mask.

Unless all hosts that are connected to the same physical network use the same subnet mask, routing problems may occur. For example, a host with IP address 9.132.56.80 and subnet mask 255.0.0.0 assumes that all hosts between 9.0.0.1 and 9.255.255.254 are attached to a local physical network, and that direct communication (without an intermediate router) is possible. The same host using subnet mask 255.255.255.0 assumes that only hosts between 9.132.56.1 and 9.132.56.254 are locally attached, and an IP router is required to transmit data to hosts outside this range.

- Hosts within the same routing domain must use the same subnet mask unless the routing protocol provides support for *variable subnet masking*.

A routing domain, see Chapter 3, “Dynamic IP Routing Protocols - Introduction” on page 91, is comprised of routers exchanging routing information through dynamic routing protocols. RIP Version 1, which is one of the dynamic routing protocols supported by the 3746 IP router, assumes that all hosts within a routing domain use the same subnet. Routing problems might result if this is not the case.

Note: OSPF Version 2 and BGP-4, which are the other routing protocols supported on the 3746 IP, do not have this problem as both provide support for variable subnet masking.

2.7.1.1 Types of Subnetting

There are two types of subnetting: static and variable length. Variable length is the more flexible of the two. Which type of subnetting is available depends upon the routing protocol being used; native IP routing supports only static subnetting, as does the widely used RIP Version 1 protocol. Chapter 3, “Dynamic IP Routing Protocols - Introduction” on page 91 discusses routing protocols in detail.

Static Subnetting: Static subnetting means that all subnets in the subnetted network use the same subnet mask. This is simple to implement and easy to maintain, but it implies wasted address space for small networks. For example, a network of four hosts that uses a subnet mask of 255.255.255.0 wastes 250 IP addresses. It also makes the network more difficult to reorganize with a new subnet mask. Currently, almost every host and router supports static subnetting.

Variable Length Subnetting: When variable length subnetting is used, the subnets that make up the network may use different subnet masks. A small subnet with only a few hosts needs a subnet mask that accommodates only these few hosts. A subnet with many hosts attached may need a different subnet mask to accommodate the large number of hosts. The ability to assign subnet masks according to the needs of the individual subnets will help conserve network addresses. Also, a subnet can be split into two parts by adding another bit to the subnet mask. Other subnets in the network are unaffected by the change. Not every host and router supports variable length subnetting.

Only networks of the size needed will be allocated and routing problems will be solved by isolating networks with routers that support variable subnetting. A host that does not support this kind of subnetting would have to route to a router that supports variable subnetting.

Mixing Static and Variable Length Subnetting: At first sight, it appears that the presence of a host that only supports static subnetting would prevent variable length subnetting from being used anywhere in the network. Fortunately this is not the case. Provided that the routers between subnets with different subnet masks are using variable length subnetting, the routing protocols employed are able to hide the difference between subnet masks from the hosts in a subnet. Hosts can continue to use basic IP routing and offload all of the complexities of the subnetting to dedicated routers.

2.7.1.2 Multiple Logical Subnets Per Physical Network

The most straightforward is to map subnet numbers one to one onto physical networks (that is, a separate subnet assigned to each physical network). There are, however, circumstances in which multiple logical subnets per physical network are quite useful. One of the more common is when it is planned that a group of workstations will be put on their own physical network but the routers to the new physical network need to be tested first. If a rule of one subnet per physical network is enforced, the addresses of the workstations must be changed every time the gateway is tested. If they may be assigned addresses using a new subnet number while they are still on the old physical network, no further address changes are needed.

By assigning multiple IP addresses to a single network interface, stations (for example, routers) can be part of multiple subnets at the same time. When multiple subnets have been defined within the same physical network, inefficient routing may result.




Figure 45. A Single Physical Network - Routing between Subnets

Figure 45 depicts an example of hosts A and B that are connected to the same physical LAN, but are part of different subnets. Router R is part of both subnets and provides routing functions between these. It would be more efficient if hosts A and B communicate directly. However, this either requires that both hosts are part of the same subnet, or a static routing statement has been added to both hosts.

2.7.2 Broadcasting

The subnet strategy requires special provisions in the network software. A host identifier of all zeroes or all ones has a special meaning. All zeroes (0, when using subnet mask 255.255.255.0) is reserved for machines that do not know their address. In certain circumstances it is possible for a machine not to know the number of the network it is on, or even its own host address. For example, 0.0.0.23 would be a machine that knew it was host number 23, but did not know on what network.

All ones (255, when using subnet mask 255.255.255.0) is used for *broadcast*. A broadcast is a message to every system on the network on which the broadcast message is sent. Broadcasting is used because it is less expensive to send a single broadcast than to send datagrams individually to each host that is interested in the information. In order to send a broadcast, you use an address that is made by using your network address, with all ones in the part of the address where the host number goes. For example, if you are on network 9.132.56, you would use 9.132.56.255 for broadcasts. How this is actually implemented depends upon the medium. It is not possible to send broadcasts on point-to-point lines. However, it is possible on a token-ring. If you use a token-ring (MAC) address with all its bits on (all ones), every machine on the token-ring is supposed to look at that datagram.

Although the official broadcast address is comprised of the network identifier and a host address of all ones, there are some other addresses that may be treated as broadcasts by certain implementations. For convenience, the standard also allows 255.255.255.255 to be used. This refers to all hosts on the local network. It is often simpler to use 255.255.255.255 instead of finding out the network number for the local network and forming a broadcast address such as 9.132.56.255. In addition, certain older implementations may use all zeroes instead of all ones to form the broadcast address. Such implementations would use 9.132.56.0 instead

of 9.132.56.255 as the broadcast address on network 9.132.56. Finally, certain older implementations may not understand about subnets. Thus they consider the network number to be 9. In that case, they will assume a broadcast address of 9.255.255.255 or 9.0.0.0.

2.7.2.1 3746 IP and Broadcasting

The 3746 IP router is capable of sending one of the following types of broadcasts:

- Limited or *local-wire* broadcast
This uses either all zeroes (0.0.0.0) or all ones (255.255.255.255) as the broadcast address.
- *Network-directed* broadcast
This uses a broadcast address comprised of the subnetted network identifier and a host identifier of either all zeros or all ones (for example, 9.132.56.0 or 9.132.56.255).

During customization (see Figure 46), either local-wire or network-directed broadcasting is selected, using either all ones or all zeros.




Figure 46. RIP - Parameters Per IP Address

The 3746 IP router will accept the following types of broadcasts:

- Limited or *local-wire* broadcast
Destination IP address 255.255.255.255 and source IP address 0.0.0.0 are accepted.
- *Subnet-directed* broadcast
A destination IP address containing the subnetted network identifier and a host identifier of either all zeros or all ones will be accepted, if the source IP address belongs to the network interface on which it has been received.
- *Network-directed* broadcast
A destination IP address containing the (non-subnetted) network number and a host number of either all zeros or all ones will be accepted, if the source IP address belongs to the network interface on which it has been received (for example, 9.255.255.255 or 9.0.0.0).

The 3746 IP router never forwards limited broadcasts. Directed broadcasts will be forwarded, if not received from the network to which it is directed, and only if the packet was not received as a link-level (for example, token-ring LAN) broadcast. Directed broadcasts are forwarded as link-level broadcasts. The 3746 contains an option to disable the forwarding of directed broadcasts.

2.7.3 Supernetting (Route Aggregation)

There is a major problem with the use of a range of class C addresses instead of a single class B address: each network must be routed separately. Standard IP routing understands only the class A, B and C network classes. Within each of these types of network, subnetting can be used to provide better granularity of the address space within each network, but there is no way to specify that multiple class C networks are actually related. The result of this is termed the *routing table explosion* problem: a class B network of 3000 hosts requires one routing table entry at each backbone router, but if the same network is addressed as a range of class C networks, it requires 16 entries.

The solution to this problem is a scheme called *Classless Inter-Domain Routing (CIDR)*. CIDR is a *proposed standard protocol* with a status of *elective*.

CIDR does not route according to the class of the network number (hence the term classless) but solely according to the high-order bits of the IP address, which are termed the *IP prefix*. Each CIDR routing entry contains a 32-bit IP address and a 32-bit network mask, which together give the length and value of the IP prefix. This can be represented as <IP_address network_mask>. For example, <194.0.0.0 254.0.0.0> represents the 7-bit IP prefix B'1100001'. CIDR handles the routing for a group of networks with a common prefix with a single routing entry. This is the reason why multiple class C network numbers assigned to a single organization have a common prefix. This process of combining multiple networks into a single entry is termed *address aggregation* or *address summarization*. It is also called *supernetting* because routing is based upon network masks that are shorter than the *natural* network mask of the IP address, in contrast to *subnetting* where the network masks are longer than the natural mask.

Unlike subnet masks, which are normally contiguous but may have a discontinuous local part, supernet masks are *always* contiguous. The 3746 IP router uses CIDR in its BGP-4 implementation (see 3.6, "Border Gateway Protocol (BGP)" on page 120).

2.7.4 3746 IP and IP Addresses

The 3746 supports class A, B, C, and D addresses. Class A, B, and C addresses are unicast addresses. Class D addresses are for multicast services (for example, by the OSPF routing protocol).

Zero, one, or multiple IP addresses can be assigned to a network interface. A subnet mask must be assigned together with each IP address. The IP addresses must be unicast addresses. IP has no fixed limit on the number of unicast addresses per interface. However, other components (for example, ARP) have a limitation of 16 IP addresses per interface.

It is not necessary to assign an IP address to the 3746 IP serial (PPP and frame relay) interfaces. However, *unnumbered* (that is, without an IP address) interfaces cannot be addressed (for example, PING and Traceroute cannot be directed to the

interface) and require OSPF as an interior routing protocol (see Chapter 3, “Dynamic IP Routing Protocols - Introduction” on page 91).

In addition to the interface addresses a *router_ID* must be assigned. Rather than specifying a specific interface, this IP address specifies the router as a whole. The *router_ID* is used for PING, Traceroute, and OSPF messages originating from the 3746 IP router. To avoid routing problems, it is recommended that you set the *router_ID* to one of the 3746 interface IP addresses.

2.8 Fragmentation and Reassembly

TCP/IP is designed for use with many different kinds of networks. TCP has the ability to *negotiate* about datagram size. When a TCP connection first opens, both ends send the maximum datagram size they can handle. The smaller of these numbers is used for the rest of the connection.

However, when two TCP endpoints are not adjacent, one or multiple IP routers will be involved in forwarding TCP data. The problem is that endpoints do not know about all the intermediate networks and may negotiate a maximum datagram size that is larger than can be accepted within one of the intermediate networks. For this reason IP routers generally implement provisions to split datagrams up into pieces. This is referred to as *fragmentation*. The IP header contains fields to indicate if a datagram has been split and information to let the pieces be put back together.

Fragmentation should, if possible, be avoided. Especially in larger networks, the extra processing required for fragmentation may lead to severe performance problems. Avoiding fragmentation on intermediate routers requires two things:

- Define the maximum packet size allowed on each of your router network interfaces.

Using the maximum packet size reduces the chance that datagrams have to be fragmented.

- Define *acceptable* maximum packet sizes on the TCP endpoints.

Unfortunately there is not a simple answer to what acceptable means in this context. The following should be considered when defining an *acceptable* value for your environment, and a balance should be found:

- It is more economical to fragment on the endpoints than on intermediate routers. Intermediate routers are shared by a high number of connections, while endpoints are generally involved in a limited number of connections. Performance problems on shared resources may impact many users, while fragmentation on the endpoints is unlikely to impact the performance.
- Define a maximum size on the endpoints that is lower than any of the maximum sizes within the intermediate networks. This will avoid intermediate fragmentation and help to maximize the performance of the network as a whole.
- When the number of intermediate hops is large, a large message traverses the network much faster when split up in small packets. However, decreasing the maximum packet too much would lead to a very inefficient use of network resources, as the network would predominantly be busy

transmitting TCP and IP header. As a rule of thumb, make sure that TCP and IP header are less than or equal to 10% of the average packet.

Note: The IP header is 20 bytes long, and the TCP header is 32 bytes or longer.

- Each IP router is limited in the number of packets it can forward. Decreasing the packet size on the TCP endpoints to minimize the network transport increases the demand for routing capacity in your network. Make sure that you have adequate routing capacity.

As mentioned earlier, the optimum maximum packet size on the endpoints depends on a number of factors. In many networks a value of 1500 bytes is used as in many implementations the lowest maximum packet size (=1500 bytes) is dictated by Ethernet networks. Therefore, using a maximum packet size of 1500 on all endpoints and a maximum packet size ≥ 1500 on all intermediate routers guarantees no fragmentation on any of the intermediate nodes.

2.8.1 3746 IP and Fragmentation

The 3746 IP router will fragment a datagram when its size exceeds the *maximum transfer unit (MTU)* on the interface on which the datagram has to be forwarded. For performance reasons, it is advised that you minimize fragmentation on intermediate routers.

Fragmentation is minimized when defining the maximum MTU value accepted on each of the 3746 IP network interfaces and defining acceptable (that is, not leading to fragmentation within intermediate networks) MTUs on the IP endstations.

Note: Defining a value ≤ 1500 on the endpoints is usually adequate. A value of 576 bytes is a *safe* value, as each implementation must support an MTU of ≥ 576 bytes.

2.9 Address Resolution Program (ARP)

There was a brief discussion earlier about what IP datagrams look like. On local area networks (LANs) IP datagrams are prefixed with a LAN header that contains the local and destination LAN addresses. The question now is how to obtain the LAN address corresponding to a given Internet address. In fact, there is a separate protocol for this, called ARP (*address resolution protocol*). (Note by the way that ARP is not an IP protocol, that is, the ARP datagrams do not have IP headers.)

Suppose you are on system 9.132.56.80 and you want to forward IP datagrams to system 9.132.56.87. Your system will first verify that 9.132.56.87 is on the same LAN, so it can talk directly. Then it will look up 9.132.56.87 in its ARP table to see if it already knows the LAN address. If so, it will complete the LAN header and send the packet.

But suppose this system is not in the ARP table. There is no way to send the packet, because you need the destination LAN address first. So it uses the ARP protocol to send an ARP request that essentially says: *I need the physical LAN address for 9.132.56.87*. Every system listens to ARP requests. When a system sees an ARP request for itself, it is required to respond. So 9.132.56.87 will see the ARP request and respond with an ARP reply saying in effect: *9.132.56.87 corresponds to X'400000002489'*. The ARP requester will save this information in

its ARP table, so future packets can be sent directly. Most systems treat the ARP table as a cache and clear entries in it if they have not been used in a certain period of time.

ARP requests must be sent as *broadcasts* as there is no way that an ARP request can be sent directly to the right system. After all, the whole reason for sending an ARP request is that you do not know the LAN address. So a LAN address of all ones is used (that is, X'FFFFFFFFFFFF'). By convention, every machine on the LAN is required to pay attention to packets with this as an address. So every system sees every ARP request. They all look to see whether the request is for their own address. If so, they respond. If not, they could just ignore it. (Some hosts will use ARP requests to update their knowledge about other hosts on the network, even if the request is not for them.) ARP is not used for packets whose IP address indicates a broadcast or multicast address, as these are always sent with the LAN broadcast address.

2.9.1 ARP and Subnets

The ARP protocol remains unchanged in the presence of subnets. Remember that each IP datagram first goes through the IP routing algorithm. This algorithm selects the hardware device driver that should send out the packet. Only then is the ARP module associated with that device driver consulted.

2.9.2 Proxy-ARP or Transparent Subnetting

Consider an IP network that is divided into subnets, interconnected by routers, but contains hosts that do not support subnetting and are unaware of the existence of multiple physical networks. For example, see hosts 1 and 2 in Figure 47. Both hosts are part of the same IP network but are attached to different physical networks. Proxy-ARP functions on A enable host 1 to learn the hardware address of 2. Routing functions on A enable IP routing between host 1 and 2. By assigning a different subnet to each LAN interface, A is capable of routing IP traffic from 1 to 2, and vice versa.




Figure 47. Hosts Interconnected by a Proxy-ARP Server/Router

When host 1 wants to send an IP datagram to host 2, it assumes that they are both attached to the same physical network and sends an ARP request to determine the physical network address of host 2. As the ARP request is sent as a LAN broadcast it will not reach host 2. However, it will reach the proxy-ARP server A.

It is the responsibility of the proxy-ARP server to return an immediate ARP reply or to return the ARP request originating from 2. Before returning the ARP reply, the proxy-ARP server will set the target protocol address field to the hardware address of the LAN to which host 1 is connected. From this moment on, host 1 thinks that

host 2 is physically attached to the same LAN, but IP traffic destined for host 2 will be sent to A instead.

To enable A to forward IP data between 1 and 2, A may (see 2.9.2.1, “Proxy-ARP - RFC 925 and RFC 1027”) assign a different subnet to each of its network interfaces. The result is transparent subnetting:

- Normal hosts (such as 1 and 2) don't know about subnetting, so they use the *old* IP routing algorithm.
- The proxy-ARP server/router between subnets has to:
 1. Use the *subnet* IP algorithm
 2. Use a modified ARP module, which can reply on behalf of other hosts




Figure 48. Proxy-ARP Server/Router

2.9.2.1 Proxy-ARP - RFC 925 and RFC 1027

The proxy-ARP concept has been introduced in *RFC 925 — Multi-LAN Address Resolution*. A subset of the RFC 925 functions are described in *RFC 1027 — Using ARP to Implement Transparent Subnet Gateways*.

RFC 925 describes a more general approach in which the proxy-ARP server maintains an ARP cache for all its LAN attachments. Each time an ARP request is received the ARP cache is checked. If an ARP request is received and the ARP cache shows that the target host is on a different LAN, the proxy-ARP server returns an ARP reply with the hardware address equal to the server interface MAC address. If an ARP request is received and no ARP entry exists for the target host, the ARP request is broadcast on all LAN interfaces (except the interface on which the ARP request has been received). When the proxy-ARP server receives the ARP reply, the target hardware address is set to the server interface MAC address on which the ARP request has been received, and the ARP reply is forwarded to the source station.

Returning the ARP reply on the proper network interface requires that the proxy-ARP server keeps a reference for each ARP request that has been forwarded but not yet responded to and the LAN interface on which the ARP request has been received. In addition, a mechanism is required to delete this reference if after a timeout period the ARP request has not been answered.

RFC 925 minimizes ARP broadcasting by relying on an ARP cache. ARP entries are learned in two ways: by caching the IP and hardware address of the station that originates the ARP request and by caching the IP and hardware address of the station that originates the ARP reply.

The concept proposed by RFC 1027 is slightly different. The proxy-ARP server ignores ARP requests if the target IP and source IP address indicate that both stations are attached to the same LAN. If not, the proxy-ARP server looks in its routing table to see on which interface the ARP request has to be forwarded. When the proxy-ARP server receives an ARP reply, the hardware address of the target station is modified and set to the interface MAC address on which the ARP reply is forwarded to the source station.

RFC 1027 requires less ARP broadcasting and the implementation is simpler. No ARP caching is required, and the timeout mechanism needed to delete unresolved ARP queries is not needed. RFC 927 forwarding decisions, both for ARP request/replies and successive data, can be fully based on the contents of the ARP caches. RFC 1027, however, requires that the proxy server has subnet awareness and contains IP routing tables to enable routing of IP data and ARP requests and replies.

2.9.3 3746 IP and ARP

The 3746 IP router supports ARP requests and responses. For details on 3746 IP ARP configuration, see 7.11, "Proxy-ARP Parameter Configuration" on page 325.

An ARP request is sent when the requested address is not found in the ARP cache table. Requests are broadcast on the local physical network. An ARP request is also sent when attempting to refresh an existing entry. In this case, the ARP request is directed to the current entry's hardware address.

A received ARP request contains the requester's protocol and hardware address and the destination IP address. ARP responds by providing the 3746 IP router's IP address to the interface on which the ARP request is received. The requester's address translation is also gleaned from the received ARP request, and an entry is made in the ARP table cache.

An ARP response is sent when a valid ARP request is received. An unsolicited ARP response is broadcast when a new local hardware IP address pair is registered, or a previously registered pair is changed. When an ARP response is received, the address translation information is cached in the ARP cache.

2.9.3.1 ARP Caching

To improve performance, ARP entries are cached in the ARP table. Each 3746 processor maintains independent ARP tables. The table contains dynamic and permanent entries. Dynamic entries are learned, while permanent entries are operator defined.

Each ARP entry contains a 6-byte token-ring LAN address and protocol-specific data referred to as *side-care*. Side-care information maintained for IP over token-ring is the *routing information field (RIF)*. On the 3746 IP router, the support for source-route bridging (SRB) is active by default.

ARP table entries that have not been recently used or updated (refreshed) are discarded. Permanent entries are never discarded. The default time for an entry to be retained before being discarded is five minutes. A configured time of '0' will retain all entries indefinitely. User commands exist to delete ARP entries. Dynamic ARP table entries are also removed when a network interface is considered down.

The 3746 IP router can be configured to attempt to refresh an unused entry before it ages out. This is referred to as the *auto-refresh* function. If configured, ARP will send out a request to the hardware address in the table to verify the address. The request is sent out approximately 30 seconds before an entry ages out. The default is to disable this function.

The ARP table sizes are limited by available memory. ARP will continue to request memory for available entries until the 3746 processor memory is depleted, at which time the IP routing function is restarted. Care should be taken in setting aging timers as it will impact the ARP table size if the aging timer is disabled.

2.9.3.2 Proxy-ARP

The 3746 IP router supports proxy-ARP. By user configuration, either RFC 925 or RFC 1027 can be selected. RFC 925 proxy-ARP is referred to as *ARP net routing*, and RFC 1027 proxy-ARP is referred to as *ARP subnet routing*. The proxy-ARP support is optional; it can be enabled or disabled for the whole 3746 IP, but it cannot be enabled or disabled on an interface basis. It is disabled by default.

Proxy-ARP is an appendage to ARP. When enabled and an ARP request is received in which the destination address does not match any of the router's IP addresses, proxy-ARP is called. Proxy-ARP looks up the destination protocol address in the routing table, and, in general, if it finds an active route to the destination, it tells ARP to send an ARP reply. No ARP reply is returned if:

- There is no route to the destination IP address in the ARP request
- The source IP address in the ARP request is not in a network directly attached to the 3746 IP router
- Source and destination IP address do not have the same network identifier

Note: An option exists to disable this.

- The physical transmit interface of the route to the destination IP address is the same interface from which the ARP request is received
- The route found for the destination protocol address is a filter (see 2.10.1, "3746 IP Filters" on page 86)

2.10 3746-9x0 Security Considerations and Controlling IP Routing

The 3746 IP router offers two methods to control its IP forwarding functions. These are:

- Filters
- Access controls

These two functions are similar in some aspects but differ in the complexity and processing needed within the 3746. IP filters are processed as a part of the normal IP routing process, but they are not as flexible as access controls. Access controls offer a lot more flexibility than filters but this is offset by the increased processing needed by access controls.

Note: For configuration details, see 5.12, "3746-9x0 Security" on page 229.

2.10.1 3746 IP Filters

The 3746 IP filters are essentially user-defined static routes. Any packet received or originated by the 3746 IP router in which the destination IP address matches a filter is discarded. No error message will be generated. Each filter, such as a static route, consists of an IP address and a mask. When an IP datagram is forwarded, the destination address is ANDed (logical AND operation) with the filter mask. If the result is the filter address, then the datagram is discarded.

```
destination_address & filter_mask = filter_address ==> discard!
```

The 3746 IP router installs filters in its routing table; therefore the number of filters is limited by the size of the routing table. Routes that match a filter are not advertised when using a dynamic routing protocol.

2.10.2 3746 Access Control

Access control provides a way for the user to control which IP packets are forwarded by the 3746 IP router based on the packet's source IP address, destination IP address, IP protocol number, and TCP or UDP port number. It is applied to all IP packets that are received or originated by the 3746 IP router. Access control is an optional function that, by default, is disabled.

Multiple access entries can be specified, each having the following parameters:

Parameter	Description
<i>Access control type</i>	The access control type can be <i>inclusive</i> or <i>exclusive</i> , indicating that if this entry is matched the datagram will be forwarded or discarded, respectively.
<i>Source IP address</i>	Value compared to the datagram's source IP address.
<i>Source IP address mask</i>	Value ANDed (logical AND operation) with the datagram's source IP address before comparing it with the access control entry's source IP address value.
<i>Destination IP address</i>	Value compared to the datagram's destination IP address.
<i>Destination IP address mask</i>	Value ANDed (logical AND operation) with the datagram's destination IP address before comparing it with the access control entry's source IP address value.
<i>First protocol number</i>	The lower bound of a range of IP protocol numbers; a datagram matches the access control entry only if its IP protocol number is greater than or equal to this value.
<i>Last protocol number</i>	The upper bound of a range of IP protocol numbers; a datagram matches the access control entry only if its IP protocol number is less than or equal to this value.
<i>First port number</i>	The lower bound of a range of TCP/UDP port numbers; a packet in which the IP protocol number indicates TCP or UDP, matches the access control entry only if its TCP or UDP port number is greater than or equal to this value.

Last port number

The upper bound of a range of TCP/UDP port number; a packet in which the IP protocol number indicates TCP or UDP, matches the access control entry only if its TCP or UDP port number is less than or equal to this value.

Access lists are ordered, meaning that if an IP datagram would match multiple entries in the list, the first matching entry in the list is the one that takes effect. If no matching access control entry is found, the 3746 discards the datagram.

Note: CCM always adds an all-inclusive non-deletable entry at the end of the list, causing all datagrams to be forwarded that do not match entries in the access control list. Make sure when using the TELNET operator interface that this entry is present and has not been deleted.

2.10.3 Rules for Filters and Access Controls

A filter on a destination IP address is added to the IP routing table with indication *filter*. The position of a route in the routing table depends on its IP address. Filters are saved into the main routing table maintained on the CBSP and the routing table cache on each of the 3746 processors.

A filter is a route similar to a static route. The primary difference is that filters are not advertised by RIP or BGP, and if a filter is replaced by another type of route (for example, OSPF) that later goes away, the filter is not automatically restored.

OSPF internal routes (intra-area and inter-area) take preference over filter routes. This is to prevent unexpected reachability problems in OSPF networks. Because it is a link-state protocol, OSPF does not take filters into account in its advertisements, so filters in an OSPF network can cause reachability problems. To be sure to filter the traffic, access control must be used instead.

Filters have another shortcoming; they can be partially overridden by a more specific route. Assume a filter for 9.0.0.0 with mask 255.0.0.0 and a route for 9.100.0.0 with mask 255.255.0.0 has been installed. Because the routing table search looks for the most specific match, datagrams to 9.100.x.x will follow the route instead of the filter.

Access control entries are searched sequentially for each datagram until a matching entry is found. This may lead to considerable overhead when many access control entries have been specified. There is no special search of a routing table for a filter. The routing table is searched once for the best matching route, and if that match turns out to be a filter, the datagram is dropped.

In general the use of filters is advised rather than using access control. However, users should be aware of the shortcomings of filters mentioned earlier.

In situations where a mixture of input, output, and global access lists, and filters have been defined, the order in which these will be applied are:

1. Global access list on 3746-9x0 input processor
2. Filter on 3746-9x0 input processor

2.11 Bootstrap Protocol (BootP)

TCP/IP provides functions to use diskless hosts as workstations, routers, terminal concentrators and so on in a TCP/IP networking environment. Diskless hosts require a mechanism to boot remotely over a network. The BootP protocol is used for remote booting over IP networks. It allows a minimum IP protocol stack with no configuration information, typically stored in ROM, to obtain enough information to begin the process of downloading the necessary boot code. BootP does not define how the downloading is done, but this process typically uses TFTP, as described in *RFC 906 — Bootstrap Loading Using TFTP*.

During the BootP process the BootP client (the diskless station) sends a BootP request to its server. When the client and server are not attached to the same subnet, a BootP relay station (for example, the 3746 IP router) can be used to provide connectivity between the client and server station. The BootP request and reply message use the format depicted in Figure 49.




Figure 49. BootP Message Format

The most important fields are:

- Code, indicating a BootP request or a reply.
- Client IP address, set by the client (either its known IP address or 0.0.0.0).
- Your IP address, the client IP address set by the server if the client IP address was 0.0.0.0.
- Server IP address, indicating BootP server.
- Router IP address, set by the 3746 IP BootP relay process.
- Client hardware address, which is the token-ring address of the client station. This address is used by the server to identify which registered client is booting.

During the BootP process the client determines its hardware address, its IP address, its server IP address, and its router address. The latter is required if the client and server are not within the same subnet. The IP addresses are optional fields in the BootP request, and the client hardware address is mandatory. If no IP addresses have been configured, the IP addresses are set to 0.0.0.0.

The BootP request is sent as a UDP datagram using port number 67. Within the IP header, the source IP address is set to the client IP address (0.0.0.0 if unknown). The destination IP address is set to the server address (255.255.255.255 if unknown). The UDP frame is sent as a link-level broadcast packet.

After the server has received the datagram (either directly or via a router) and looks up the hardware address of the client in its configuration file, fills in the remaining fields in the UDP datagram and returns a BootP reply using UDP port 68.

The BootP reply can simply be returned if the client IP address was included in the BootP request. ARP processing is used to learn the hardware address corresponding to the client IP address.

If the client IP address is learned from the server configuration, ARP cannot be used to find the client hardware address as the client does not know its IP address and, therefore, will not reply to an ARP request. To this *chicken and egg* problem are two possible solutions:

1. If the server has a mechanism for directly updating its own ARP cache without using ARP itself, it does so and then sends the datagram directly.
2. If the server cannot update its own ARP cache, it must send a broadcast reply.

When the client receives the reply, the BootP client will record its own IP address (allowing it to respond to ARP requests) and begin the bootstrap process.

2.11.1 3746 IP and BootP

One restriction with the scheme described in the previous section is related to the use of the limited broadcast address for BootP requests; it requires that the server is on the same subnet as the recipient.

The 3746 IP router, however, has implemented a mechanism to forward BootP requests to one or multiple preconfigured BootP servers (for configuration details, see 7.7, “IP Parameter Configuration” on page 307). If the BootP request's router address is zero, the 3746 BootP relay agent sets it to the IP address of its receiving interface.

The 3746 IP BootP relay agent discards the BootP request message when:

- The BootP request's hops value is greater than the value configured by the user.
- The BootP request's hops value is greater than or equal to 16.
- The BootP request's seconds (retry) value is greater than the value configured by the user.

This allows BootP servers on the local network to respond before the BootP relay agents forward to remote BootP servers.

The BootP relay agent forwards a BootP reply message to the interface whose IP address matches the BootP reply's router IP address. Because the BootP client may not know its IP address and, therefore, may not be able to respond to an ARP request, the 3746 IP BootP relay agent installs an entry in the outgoing interface ARP cache table using the client IP and hardware address specified in the BootP reply message.

Once a BootP server has responded and the BootP client has processed the reply, it may proceed with the transfer of the boot file and execute the full boot process. The transfer of the boot file is usually done using the *trivial file transfer protocol* (TFTP). The full boot process will replace the minimum IP protocol stack used by BootP and TFTP by a normal IP protocol stack transferred as part of the boot file and containing the correct customization for the client.

2.12 Request for Comments (RFCs)

Publications like the one you are reading can only give limited information on the many protocols involved in TCP/IP. Fortunately, all official internet protocols have been standardized and are documented in RFCs (*request for comments*).

If you require more detailed information than is given in this publication, we refer you to the official Internet publications. RFCs can be retrieved from many sources. For example, RFCs can be obtained in printed form from:

Network Solutions, Inc.
Attn: InterNIC Registration Service
505 Huntmar Park Drive
Herndon, VA 22070

Help Desk Telephone Number:
703-742-4777

FAX Number 703-742-4811

<http://www.internic.net>

To get RFCs electronically, users may use anonymous FTP to ds.internic.net (198.49.45.10) and retrieve files from the directory rfc, or Gopher to internic.net (198.41.0.5).

A good starting point to find the RFC you are looking for is RFC 1780 - *Internet Official Protocol Standards* and RFC 1700 - *Assigned Internet Number*.

Chapter 3. Dynamic IP Routing Protocols - Introduction

The IBM 3746-9x0 supports a number of IP routing (routing table maintenance) protocols. These allow the 3746 to exchange IP routing information with other 3746s and with other similarly capable IP routers. The following are supported:

- RIP Version 1
- OSPF Version 2
- BGP Version 4

RIP and OSPF are referred to as *interior gateway¹ protocols* (IGPs), while BGP is referred to as an *exterior gateway protocol* (EGP).

In the following sections the above routing protocols are discussed. For 3746 IP configuration details, see Chapter 4, "IBM 3746-9x0 and IP Routing Protocols" on page 131.

3.1 Interior and Exterior Gateway Protocols




Figure 50. Autonomous Systems

Gateway protocols are referred to as interior or exterior depending on whether they are used within or between Autonomous Systems (ASs). allow routers to exchange routing information within an AS. Exterior gateway protocols allow the exchange of summary reachability information between separately administered ASs.

ASs are defined as logical portions of larger IP networks that are administered by single authorities. An AS would normally comprise the internetwork within an organization and would be designated as such to allow communication over public

¹ The term gateway is being used as is the convention for TCP/IP specialists (that is, to mean router). It is not being used in the strict ISO OSI sense.

IP networks with ASs belonging to other organizations. It is mandatory to register an organization's internetwork as an AS in order to use these public IP services.

Figure 50 on page 91 illustrates two ASs interconnected by routers. It shows how IGPs are used within the ASs and an EGP between them.

ASs must be registered publicly. If you require an AS number, or an IP network address, to allow your network to connect to public IP services you should contact:

Network Information Center
GSI
14200 Park Meadow Drive, Suite 200
Chantilly, VA 22021
USA
Tel: +1 703 802 4535

3.2 Choosing Gateway Protocols

Within an AS (or if you are building a private IP network) you are free to choose the interior gateway protocol, or combination of protocols, that best meets your needs.

Each interior gateway protocol, however, has different characteristics and selection must be carried out carefully to meet internetwork design requirements. The design considerations for selecting interior gateway protocols for an IP network using 3746 IP routers are described in 4.1, "Using RIP" on page 131 and in 4.2, "Using OSPF" on page 134.

If you wish to communicate with other ASs you are once again, in principle, free to choose the exterior gateway protocol that best meets your needs. The interior gateway protocol used within an AS is not constrained by the choice of an exterior gateway protocol. However, there is synergy between some interior and exterior gateway protocols (for example, OSPF and BGP).

In practice AS-to-AS communication is governed by rules set by the administrators of the public Internet or by private IP service providers. An internetwork design must accommodate the exterior gateway protocols required by the IP service provider.

The design considerations for the exterior gateway protocol supported by the 3746 IP router are described in 4.3, "Using BGP Version 4" on page 140.

3.3 Routing Algorithms

Interior and exterior gateway protocols currently implemented in the IBM 3746-9x0 use one of two generic classes of dynamic routing algorithm known as *distance vector* and *link-state* routing algorithms.

Dynamic routing algorithms allow routers to exchange route or link information, from which the best paths to reachable destinations in an internetwork are calculated.

Static routing may also be used to supplement dynamic routing.

3.3.1 Static Routing

Static routing requires that routes be configured manually into a router.

Normally manual configuration is to be avoided, particularly within an AS, but there are circumstances when static routing can be attractive:

- To define a default route, or a route that is not being advertised within a network
- To supplement or replace exterior gateway protocols:
 - When line tariffs between ASs are high or based on traffic volumes it may be desirable to avoid the cost of routing protocol traffic
 - If complex routing policies are to be implemented
 - To avoid disruption caused by faulty exterior gateways in other Autonomous Systems

3.3.2 Distance Vector Routing

Distance vector routing is still the most common type of routing encountered in IP internetworks.

The principle behind distance vector routing is very simple. Each router in an internetwork maintains the distance from itself to every known destination in a *distance vector table*. Distance vector tables consist of a series of destinations (vectors) and costs (distances) to reach them and least cost route to destinations to be determined at the time of transmission.

The distances in the tables are computed from information provided by neighbor routers that transmit their own distance vector tables across shared networks. The sequence of operations for doing this is as follows:

- Each router is configured with an identifier and the cost of each of its network links (the cost is normally fixed at 1 and hence equal to a single hop, but can reflect some other measurement taken for the link).
- Each router initializes with a distance vector table containing zero for itself, one for directly attached networks and infinity for every other destination.
- Each router periodically (typically every 30 seconds) transmits its distance vector table to each of its neighbors (it may also transmit it when a link first comes up or when the table changes).
- Each router saves the most recent table it receives from each neighbor and uses the information to calculate its own distance vector table.
- The total cost to each destination is calculated by adding the cost reported to it in a neighbor's distance vector table to the cost of the link to that neighbor.
- The distance vector table (the routing table) for the router is then created by taking the lowest cost calculated for each destination.

Figure 51 on page 94 shows the distance vector tables for three routers within a simple internetwork.




Figure 51. Distance Vector - Routing Table Calculation

The distance vector algorithm produces a stable routing table after a period directly related to the number of router hops across the network. This period is referred to as the *convergence time* for a network and represents the time it takes for distance vector information to traverse the network. It is possible in large internetworks for this time to be too long to be useful.

Routing tables are recalculated if a changed distance vector table is received from a neighbor, or if a link to a neighbor is found to have gone down. If a network link goes down, the distance vector tables that have been received over it are discarded and the routing table recalculated.

The chief advantage of a distance vector is that it is very easy to implement. There are also significant disadvantages, the main ones being the instability caused by old routes persisting in an internetwork, the long convergence time on large internetworks, the limit to the size of an internetwork imposed by maximum hop counts, and the fact that distance vector tables are always transmitted even if their contents have not changed.

Enhancements to the basic algorithm have evolved to overcome the first two of these problems. They are described in the following subsections.

3.3.2.1 Split Horizon

The basic distance vector algorithm will always allow a router to correctly calculate its distance vector table.

Using the example shown in Figure 52 you can see one of the problems of distance vector protocols known as *counting to infinity*.




Figure 52. Counting to Infinity - Example Network

Counting to infinity occurs when a network becomes unreachable, but erroneous routes to that network persist because of the time for the distance vector tables to converge.

The example network shows four routers interconnected by five network links. The networks all have a cost of one except for that from C to D which has a cost of 10.

Each of the routers A, B, C and D have routes to all networks. If we show only the routes to the target network we will see they are as follows:

For D : Directly connected. Metric 1

For B : Route via D. Metric 2

For C : Route via B. Metric 3

For A : Route via B. Metric 3

If the link from B to D fails, then all routes will adjust in time to use the link from C to D. The convergence time for this, however, may be considerable.

Distance vector tables begin to change when B notices that the route to D has become unavailable. Figure 53 on page 96 shows how the routes to the target network will change, assuming all routers send distance vector table updates at the same time.

Time													
D:	Direct	1	Direct	1	Direct	1	Direct	1	----	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	----	C	11	C	12

Figure 53. Counting to Infinity

The problem can be seen clearly. B is able to remove the failed route immediately because it times out the link. Other routers, however, have tables that contain references to this route for many update periods after the link has failed.

1. Initially A and C have a route to D via B.
2. Link from D to B fails.
3. A and C then send updates based on the route to D via B even after the link has failed.
4. B then believes it has a route to D via either A or C. It has not in reality, as the routes are vestiges of the route via B, which has failed.
5. A and C then see that the route via B has failed, but believe a route exists via one another.

Slowly the distance vector tables converge, but not until the metrics have counted up, in theory, to infinity. To avoid this happening practical implementations of distance vector have a low value for infinity; for example, RIP uses a maximum metric of 16.

The manner in which the metrics increment to infinity gives rise to the term *counting to infinity*. It occurs because A and C are engaged in an extended period of mutual deception, each claiming to be able to get to the target network D via one another.

Counting to infinity can easily be prevented if a route to a destination is never reported in the distance vector table sent to the neighbor from which the route was learned. *Split horizon* is the term used for this technique.

The incorporation of split horizon would modify the sequence of distance vector table changes to that shown in Figure 54 on page 97. The tables can be seen to converge considerably faster than in Figure 53.

Time								
D:	Direct	1	Direct	1	Direct	1	Direct	1
B:	Unreachable		Unreachable		Unreachable		C	12
C:	B	3	A	4	D	11	D	11
A:	B	3	C	4	Unreachable		C	12
Note: Faster routing table convergence								

Figure 54. Split Horizon

3.3.2.2 Split Horizon with Poison Reverse

Poison reverse is an enhancement to split horizon whereby routes learned from a neighbor router are reported back to it but with a metric of infinity (that is, network unreachable).

The use of poison reverse is safer than split horizon alone because it breaks erroneous looping routes immediately.

If two routers receive routes pointing at each other, and they are advertised with a metric of infinity, the routes will be eliminated immediately as unreachable. If the routes are not advertised in this way they must be eliminated by the timeout that results from a route not being reported by a neighbor router for several periods (for example, six periods for RIP).

Poison reverse does have one disadvantage. It significantly increases the size of distance vector tables that must be exchanged between neighbor routers because all routes are included in the distance vector tables. While this is generally not a problem on LANs, it can cause real problems on point-to-point connections in large internetworks.

3.3.2.3 Triggered Updates

Split horizon with poison reverse will break routing loops involving two routers.

It is still possible, however, for there to be routing loops involving three or more routers. For example, A may believe it has a route through B, B through C and C through A. This loop can only be eliminated by the timeout that results from counting to infinity.

Triggered updates are designed to reduce the convergence time for routing tables and, hence, reduce the period during which such erroneous loops are present in a internetwork.

When a router changes the cost for a route in its distance vector table it must send the modified table immediately to neighbor routers. This simple mechanism ensures that topology changes in a network are propagated quickly, rather than at a rate dependent on normal periodic updates.

3.3.3 Link-State Routing

The growth in the size of internetworks in recent years has necessitated the replacement of distance vector routing algorithms with alternatives that address the shortcomings identified in 3.3.2, “Distance Vector Routing” on page 93.

These new protocols have been based on *link-state* or shortest path first algorithms. The best example is the OSPF interior gateway protocol.

The principle behind link-state routing is straightforward, although implementation can be complex:

- Routers are responsible for contacting neighbors and learning their identities.
- Routers construct link-state packets that contain lists of network links and their associated costs.
- Link-state packets are transmitted to all routers in a network.
- All routers, therefore, have an identical list of links in a network and can construct identical topology maps.
- The maps are used to compute the best routes to all destinations.

Routers contact neighbors by sending *Hello* packets on their network interfaces. Hello packets are sent directly to neighbors on point-to-point and non-broadcast networks. On LAN networks, Hello packets are sent to a predefined group or multicast IP address that can be received by all routers. Neighbors who receive Hellos from a router should reply with Hello packets that include the identity of that originating router.

Once neighbors have been contacted in this way, link-state information can be exchanged.

Link-state information is sent in the form of *link-state information packets* (LSPs), also known as link-state advertisements. LSPs provide the database from which network topology maps can be calculated at each router. LSPs are normally sent only under the following specific circumstances:

- When a router discovers a new neighbor
- When a link to a neighbor goes down
- When the cost of a link changes

Once a router has generated an LSP, it is critical that it is received successfully by all other routers in a network. If this does not happen, routers on the network will calculate network topology based on incorrect link-state information.

Distribution of LSPs would normally be on the basis of each router's routing tables. However, this leads to a *chicken and egg* situation. Routing tables would rely on LSPs for their creation and LSPs would rely on routing tables for their distribution. A simple scheme called *flooding* overcomes this and ensures that LSPs are successfully distributed to all routers in a network.

Flooding requires that a router that receives an LSP must transmit it to all neighbors except the one from which it was received. All LSPs must be explicitly acknowledged to ensure successful delivery, and they are sequenced and time stamped to ensure duplicates are not received and retransmitted.

When a router receives an LSP it looks in its database to see the sequence number of the last LSP from the originator. If the sequence number is the same as, or earlier than, the sequence number of the LSP in its database, then the LSP is discarded. Otherwise the LSP is added to the database.

The flooding process ensures that all routers in a network have the same link-state information. All routers are then able to compute the same shortest path tree topology map for the network and, hence, select best routes to all destinations.

3.4 Routing Information Protocol (RIP)

The *routing information protocol* (RIP) is an interior gateway protocol defined in RFC 1058.

It is an IAB standard protocol; its status is elective. This means that it is one of several interior gateway protocols available, and it may or may not be implemented on a system. If a system does implement it, however, the implementation should be in line with the RFC.

RIP is based on the Xerox PUP and XNS routing protocols. The RFC was issued after many RIP implementations had been completed. For that reason some do not include all the enhancements to the basic distance vector routing protocol (such as poison reverse and triggered updates).

RIP is very widely used because the code (known as *ROUTED*) was incorporated on the Berkeley Software Distribution (BSD) UNIX operating system and in other UNIX systems based on it.

The next sections provide an overview of the RIP protocol. The IBM 3746-9x0 supports a full implementation of RIP Version 1. For details on the 3746 IP implementation, see 3.4.2, “3746 IP and RIP” on page 101.

3.4.1 Protocol Description

RIP is a standard distance vector routing protocol, as described in 3.3.2, “Distance Vector Routing” on page 93.

RIP packets are transmitted onto a network in *User Datagram Protocol* (UDP) datagrams, which, in turn, are carried in IP datagrams. RIP sends and receives datagrams using UDP port 520. RIP datagrams have a maximum size of 512 bytes and tables larger than this must be sent in multiple UDP datagrams.

RIP datagrams are normally broadcast onto LANs using the LAN MAC All-Stations broadcast address and the IP network or subnetwork broadcast address. They are specifically addressed on point-to-point and multi-access non-broadcast networks, using the destination router IP address.

Routers normally run RIP in *active mode*, that is, advertising their own distance vector tables and updating them based on advertisements from neighbors. End nodes, if they run RIP, normally operate in *passive (or silent) mode*, that is, updating their distance vector tables on the basis of advertisements from neighbors, but not advertising them.

RIP specifies two packet types: *request* and *response*.

A request packet is sent by routers to ask neighbors to send part of their distance vector table (if the packet contains destinations), or all their table (if no destinations are specified).

A response packet is sent by routers to advertise their distance vector table in the following circumstances:

- Every 30 seconds
- In response to a request packet
- When distance vector tables change (if triggered updates are supported)




Figure 55. Generalized RIP Packet Format

Active and passive systems listen for all response packets and update their distance vector tables accordingly. A route to a destination, computed from a neighbor's distance vector table, is kept until an alternate route is found with lower cost, or it is not re-advertised in six consecutive RIP responses. In this case the route is timed out and deleted.

The RFC defines a packet format that can be used with different network protocols. It does this by specifying an *address family identifier* that defines the type (and hence the length) of the network address. The generalized format of RIP packets is shown in Figure 55.

RIP may be used, therefore, for protocols other than IP simply by setting the address family identifier. The RFC requires that RIP response handling discards entries for unsupported address families, but processes entries for supported address families in the normal way.

When RIP is used with IP the address family identifier is 2, and the address fields are 4 bytes. To reduce problems of counting to infinity the maximum metric is 16 (unreachable) and directly connected networks are defined as having a metric of 1.

The RIP packet format for IP is shown in Figure 56 on page 101.




Figure 56. IP Specific RIP Packet Format

RIP makes no provision for passing subnet masks with its distance vector tables. A router receiving a RIP response must already have subnet mask information to allow it to interpret the network identifier and host identifier portions of the IP address correctly.

In the absence of subnet mask information a router will interpret routes as best it can. If it knows an IP network has a specific subnet mask it will interpret all other route information for that network on the basis of that single mask. If it receives a packet with bits set in the field that it regards as the host field, it will interpret it as a route to a host with a mask of 255.255.255.255.

The above makes it impossible for RIP to be used in an internetwork with variable length subnet masks.

3.4.2 3746 IP and RIP

The IBM 3746-9x0 supports RIP, this includes support for split horizon and poison reverse. The two are mutually exclusive. One of the two is always in effect. To speed network convergence, triggered updates are sent.

3.4.2.1 Route Acceptance Policy

Per interface the user can configure whether or not:

- RIP will listen
- Network, subnet, and/or host routes are learned
- Default and/or static routes can be overridden with a learned route

Note: Default routes are indicated by a destination and mask of 0.0.0.0.

3.4.2.2 Route Advertisement Policy

Per interface the user can configure whether or not:

- RIP will advertise
- Network, subnet, and/or host routes are advertised
- Default and/or static routes are advertised

Subnets will be advertised if they are part of the same natural network (that is, class A, B, or C) as the interface's IP address, and the route subnet mask must be

the same as the interface's subnet mask. Host routes (mask 255.255.255.255) are exempt from these rules.

For details on configuring 3746 IP RIP functions, see 4.1.1, “RIP Configuration on the IBM 3746-9x0” on page 132.

3.5 Open Shortest Path First (OSPF)

The *open shortest path first* (OSPF) V2 protocol is an interior gateway protocol defined in RFC 1583. A report on the use of OSPF V2 is contained in RFC 1246 - *Experience with the OSPF Protocol*.

It is an IAB standard protocol; its status is elective.

OSPF is important because it has a number of features not found in other interior gateway protocols. Support for these additional features makes OSPF the preferred choice for new IP internetwork implementations:

- Variable length subnet masks
- Alternate routes based on IP type of service (TOS)
- Equal cost multipath routes
- Faster network convergence in the presence of change, no counting to infinity
- Better scaling capability
- Use of local multicast to limit traffic to OSPF routers

The next sections provide an overview of the OSPF protocol. The IBM 3746-9x0 supports a full implementation of OSPF V2. For details on the 3746 IP implementation, see 3.5.4, “3746 IP and OSPF” on page 118.

3.5.1 OSPF Terminology

OSPF uses a specific terminology that must be understood before the protocol can be described.

3.5.1.1 Areas

OSPF internetworks are organized into *areas*.

An OSPF area consists of a number of networks and routers that are logically grouped together. Areas may be defined on a per location or a per region basis, or they may be based on administrative boundaries.

All OSPF networks consist of at least one area, the backbone, plus as many additional areas as are demanded by network topology and other design criteria.

Within an OSPF area all routers maintain the same topology database, exchanging link-state information to maintain their synchronization. This ensures that all routers calculate the same network map for the area.

Information about networks outside an area is summarized by *area border* or *AS boundary routers* (see 3.5.1.3, “Intra-Area, Area Border and AS Boundary Routers” on page 104) and flooded into the area. Routers within an area have no knowledge of the topology of networks outside the area, only of routes to destinations provided by area border and AS boundary routers.

The importance of the area concept is that it limits the size of the topology database that must be held by routers. This has a direct impact on the processing to be carried out by each router, and on the amount of link-state information that must be flooded into individual networks.

3.5.1.2 The OSPF Backbone

All OSPF networks must contain at least one area, the *backbone*, which is assigned an area identifier of *0.0.0.0*.

The backbone has all the properties of an area, but has the additional responsibility of distributing routing information between areas attached to it.

When routing a packet between two areas, the backbone is used. The path that the packet will travel can be broken up into three contiguous pieces: an intra-area path from the source to an area router, a backbone path between source and destination area, and then another intra-area path to the destination. The OSPF algorithm finds a set of such paths that have the smallest cost.

Normally an OSPF backbone should be contiguous, that is, with all backbone routers attached to one another. This may not be possible because of network topology, in which case backbone continuity must be maintained by the use of *virtual links*.

Virtual links are backbone router-to-backbone router connections that traverse a non-backbone area.

Routers within the backbone operate identically to other intra-area routers and maintain full topology databases for the backbone area.




Figure 57. OSPF Network

3.5.1.3 Intra-Area, Area Border and AS Boundary Routers

There are three possible types of routers in an OSPF network.

Routers that are situated entirely within an OSPF area are called *intra-area routers*. All intra-area routers flood router links advertisements into the area to define the links they are attached to. If elected as designated or backup designated routers (see 3.5.1.6, “Designated and Backup Designated Router” on page 105), they also flood network links advertisements to define the identity of all routers attached to the network. Intra-area routers maintain a topology database for the area in which they are situated.

Routers that connect two or more areas are referred to as *area border routers*. Area border routers maintain topology databases for each area to which they are attached, and exchange link-state information with other routers in those areas. Area border routers also flood summary link-state advertisements into each area to inform them of inter-area routes.

Routers that are situated at the periphery of an OSPF internetwork and exchange reachability information with routers in other ASs using exterior gateway protocols are called *AS boundary routers*. Routers that import static routes or routes from other IGPs such as RIP into an OSPF network are also AS boundary routers. AS boundary routers are responsible for flooding AS external link-state advertisements into all areas within the AS to inform them of external routes.

Figure 57 on page 103 shows the location of intra-area, area border and AS boundary routers within an OSPF internetwork.

3.5.1.4 Neighbor Router

Two routers that have interfaces to a common network are said to be *neighbors*.

Neighbor routers are discovered by the OSPF Hello protocol, which is described in 3.5.2.1, “Discovering Neighbors - The OSPF Hello Protocol” on page 108.

3.5.1.5 Adjacent Router

Neighbor routers may become *adjacent*. They are said to be adjacent when they have synchronized their topology databases through the exchange of link-state information.

Link-state information is exchanged only between adjacent routers, not between neighbor routers.

Not all neighbor routers become adjacent. Neighbors on point-to-point links do so, but on multi-access networks adjacencies are only formed between individual routers and the designated and backup designated routers.

The exchange of link-state information between neighbors can create significant amounts of network traffic. Limiting the number of adjacencies on multi-access networks in this way achieves considerable reductions in network traffic.

3.5.1.6 Designated and Backup Designated Router

All multi-access networks have a *designated* and a *backup designated* router.

These routers are elected automatically for each network once neighbor routers have been discovered by the Hello protocol.

The designated router performs two key roles for a network:

- It generates network links advertisements that list the routers attached to a multi-access network.
- It forms adjacencies with all routers on a multi-access network and, therefore, becomes the focal point for forwarding of all link-state advertisements.

The backup designated router forms the same adjacencies as the designated router. It therefore has the same topology database and is able to assume designated router functions should it detect that the designated router has failed.

3.5.1.7 Stub Areas

In some ASs, the majority of the topological database may contain routes to destinations outside the AS. OSPF allows certain areas to be configured as *stub areas*. OSPF external advertisements are not flooded into/throughout stub areas; routing to AS external destinations in these areas is based on default routing. This reduces the memory requirements for stub area's internal routers.

On one or more of the stub area border routers, a default route must be defined, which is then advertised into the stub area.

There are a couple of restrictions on the use of stub areas. Virtual links cannot be configured through stub areas. In addition, AS boundary routers cannot be placed internal to stub areas.

3.5.1.8 Point-to-Point and Multi-Access Networks

All OSPF areas consist of aggregates of networks linked by routers. OSPF categorizes networks into two different types.

Point-to-point networks directly link two routers. OSPF packets on a point-to-point network are multicast to the neighbor router. *Multicasting* is the term used for transmitting IP datagrams to a functional rather than a specific IP address. A functional address will typically be recognized by a number of systems and can be considered a form of limited broadcast. OSPF defines the use of two multicast addresses (224.0.0.4 and 224.0.0.5) for OSPF router interactions.

Multi-access networks are those that support the attachment of more than two routers. They are further subdivided into two types:

- Broadcast
- Non-broadcast

Broadcast networks have the capability of directing OSPF packets to all attached routers, using an address that is recognized by all of them. A token-ring LAN is an example of a broadcast multi-access network.

Non-broadcast networks do not have this capability and all packets must be specifically addressed to routers on the network. This requires that routers on a

non-broadcast network know the addresses of neighbors. A frame relay network is an example of a non-broadcast multi-access network.

3.5.1.9 Link-State Advertisements

Link-state information is exchanged by adjacent OSPF routers to allow area topology databases to be maintained, and inter-area and inter-AS routes to be advertised.

Link-state information consists of five types of *link-state advertisement*. Together these provide all the information needed to describe an OSPF network and:

1. Router links
2. Network links
3. Summary links (Type 3 and 4)
4. AS external links

Router links advertisements are generated by all OSPF routers and describe the state of the router's interfaces (links) within the area. They are flooded throughout a single area only.

Network links advertisements are generated by the designated router on a multi-access network and list the routers connected to the network. They are flooded throughout a single area only.

Summary links advertisements are generated by area border routers. There are two types; one describes routes to destinations in other areas, and the other routes to AS boundary routers. They are flooded throughout a single area only.

AS external links advertisements are generated by AS boundary routers and describe routes to destinations external to the OSPF network. They are flooded throughout all areas in the OSPF network, except stub areas.

3.5.2 Protocol Description

The OSPF protocol is an implementation of a *link-state* routing protocol, as described in 3.3.3, "Link-State Routing" on page 98.

OSPF packets are transmitted directly in IP datagrams. IP datagrams containing OSPF packets can be distinguished by their use of *protocol identifier 89* in the IP header. Therefore, OSPF packets are not contained in TCP or UDP headers. OSPF packets are always sent with IP *type of service* set to 0, and the IP *precedence field* set to internetwork control. This is to aid them in getting preference over normal IP traffic.

Further details of IP protocol identifiers, type of service and precedence can be found in RFC 791 - *Internet Protocol*.

The IP destination address for OSPF packets is selected as follows. On physical point-to-point networks, the IP destination is always set to the address AllSPFRouters (that is, 224.0.0.5). On all other network types (including virtual links), the majority of OSPF packets are sent as unicasts. In this case, the IP destination is just the neighbor IP address. The only packets not sent as unicasts are on broadcast networks; on these networks Hello packets are sent to the multicast destination AllSPFRouters. The designated router and its backup send

both Link-State Update Packets and Link-State Acknowledgment Packets to the multicast address AllSPFRouters, while all other routers send both their Link-State Update and Link-State Acknowledgment packets to the multicast address AllIDRouters (that is, 224.0.0.4).

Retransmissions of Link-State Update packets are always sent as unicasts.

The IP source address is the IP address of the sending interface. As for unnumbered point-to-point, no IP address is defined, and another IP address belonging to the router is used. On the 3746, the *router_ID* is used.

All OSPF packets share a common header, which is shown in Figure 58.

This header provides general information, such as area identifier and originating router identifier, and also includes a checksum and authentication information. A type field defines each OSPF packet as one of five possible types:

1. Hello
2. Database Description
3. Link-State Request
4. Link-State Update
5. Link-State Acknowledgement

The router identifier, area identifier, and authentication information are configurable for each OSPF router.




Figure 58. OSPF Common Header

The OSPF protocol defines a number of stages that must be executed by individual routers. They are as follows:

- Discovering neighbors
- Electing the designated router
- Initializing neighbors
- Propagating link-state information
- Calculating routing tables

The use of the five OSPF packet types to implement stages of the OSPF protocol are described in the following subsections.

During OSPF operation a router cycles each of its interfaces through a number of *states* from *Down*, through *Waiting*, to *DR Other*, *BackupDR* or *DR* (DR stands for *designated router*) depending on the status of each attached network and the identity of the designated router elected for each of them. A detailed description of these states is outside the scope of this document, but can be found in RFC 1583.

At the same time a router cycles each neighbor interface (interaction) through a number of states as it discovers them and then becomes adjacent. These states are *Down*, *Attempt*, *Init*, *2-Way*, *ExStart*, *Exchange*, *Loading* and *Full*. Once again a description of these is outside the scope of this document but can be found in RFC 1583.

3.5.2.1 Discovering Neighbors - The OSPF Hello Protocol

The Hello protocol is responsible for discovering neighbor routers on a network and establishing and maintaining relationships with them.

Hello packets are sent out periodically on all router interfaces. The format of these is shown in Figure 59.




Figure 59. OSPF Hello Packet

Hello packets contain the identities of neighbor routers whose Hello packets have already been received over a specific interface. They also contain the *network mask*, *router priority*, *designated router identifier* and *backup designated router identifier*. The final three parameters are used to elect the designated router on multi-access networks.

The network mask, router priority, Hello interval and router dead interval are configurable for each interface on an OSPF router.

A router interface changes state from *Down* to *Point-to-Point* (if the network is point-to-point), to *DR Other* (if the router is ineligible to become designated router), or otherwise to *Waiting* as soon as Hello packets are sent over it.

A router receives Hello packets from neighbor routers via its network interfaces. When this happens the neighbor interface state changes from *Down* to *Init*. Bidirectional communication is established between neighbors when a router sees itself listed in a Hello packet received from another router. Only at this point are the two routers defined as true neighbors, and the neighbor interface changes state from *Init* to *2-Way*.

3.5.2.2 Electing the Designated Router

All multi-access networks have a designated router. There is also a backup designated router that takes over in the event that the designated router fails.

The use of a backup, which maintains an identical set of adjacencies and an identical topology database to the designated router, ensures there is no extended loss of routing capability if the designated router fails.

The designated router performs two major functions on a network:

- It originates network links advertisements on behalf of the network.
- It establishes adjacencies with all other routers on the network. Only routers with adjacencies exchange link-state information and synchronize their databases.

The designated router and backup designated router are elected *router identifier*, *router priority*, *designated router* and *backup designated router* fields in Hello packets. Router priority is a single-byte field that defines the priority of a router on a network. The lower the value of the priority field the more likely the router is to become the designated router, hence the higher its priority. A zero value means the router is ineligible to become designated or backup designated router.

The process of designated router election is as follows:

1. The current values for designated router and backup designated router on the network are initialized to 0.0.0.0.
2. The current values for router identifier, router priority, designated router and backup designated router in Hello packets from neighbor routers are noted. Local router values are included.
3. Backup Designated Router Election:

Routers that have been declared as designated routers are ineligible to become backup designated routers.

The backup designated router will be declared to be:

- The highest priority router that has been declared as the backup designated router
- The highest priority router if no backup designated router has been declared

If equal priority routers are eligible, the one with the highest router identifier is chosen.

4. Designated Router Election:

The designated router will be declared to be:

- The highest priority router that has been declared the designated router
 - The highest priority router if no designated router has been declared
5. If the router carrying out the determination is declared the designated or backup designated router, then the previous steps are re-executed. This ensures that no router can declare itself both designated and backup designated router.

Once designated and backup designated routers have been elected for a network, they proceed to establish adjacencies with all routers on the network.

Completion of the election process for a network causes the router interface to change state from *Waiting* to *DR*, *BackupDR*, or *DR Other* depending on whether the router is elected the designated router, the backup designated router or neither of these.

3.5.2.3 Establishing Adjacencies - Database Exchange

A router establishes adjacencies with a subset of neighbor routers on a network.

Routers connected by point-to-point networks and virtual links always become *adjacent*. Routers on multi-access networks form adjacencies with the designated and backup designated routers only.

Link-state information flows only between adjacent routers. Before this can happen it is necessary for them to have the same topological database and to be synchronized.

This is achieved in OSPF by a process called *database exchange*.

Database exchange between two neighboring routers occurs as soon as they attempt to bring up an adjacency. It consists of the exchange of a number of database description packets that define the set of link-state information present in the database of each router. The link-state information in the database is defined by the list of link-state headers for all link-state advertisement in the database (see Figure 64 on page 113 for information on the link-state header).

The format of database description packets is shown in Figure 60.




Figure 60. OSPF Database Description Packet

During the database exchange process the routers form a *master/slave* relationship, with the master being the first to transmit. The master sends database description packets to the slave to describe its database of link-state information. Each packet is identified by a sequence number and contains a list of the link-state headers in the master's database. The slave acknowledges each packet by sequence number and includes its own database of headers in the acknowledgements.

Flags in database description packets indicate whether they are from a master or slave (the *M/S* bit), the first such packet (the *I* bit) and if there are more packets to come (the *M* bit). Database exchange is complete when a router receives a database description packet from its neighbor with the M bit off.

During database exchange each router makes a list of the link-state advertisements for which the adjacent neighbor has a more up-to-date instance (all advertisements are sequenced and time stamped). Once the process is complete, each router requests these more up-to-date instances of advertisements using link-state requests.

The format of link-state request packets is shown in Figure 61.




Figure 61. OSPF Link-State Request Packet

The database exchange process sequences the neighbor interface state from 2-Way through:

- ExStart* as the adjacency is created and the master agreed upon
- Exchange* as the topology databases are being described
- Loading* as the link-state requests are being sent and responded to
- Full* when the neighbors are fully adjacent

In this way the two routers synchronize their topology databases and are able to calculate identical network maps for their OSPF area.

3.5.2.4 Link-State Propagation

Information about the topology of an OSPF network is passed from router to router in link-state advertisements.

Link-state advertisements pass between adjacent routers in the form of *link-state update* packets, the format of which is shown in Figure 62 on page 112.

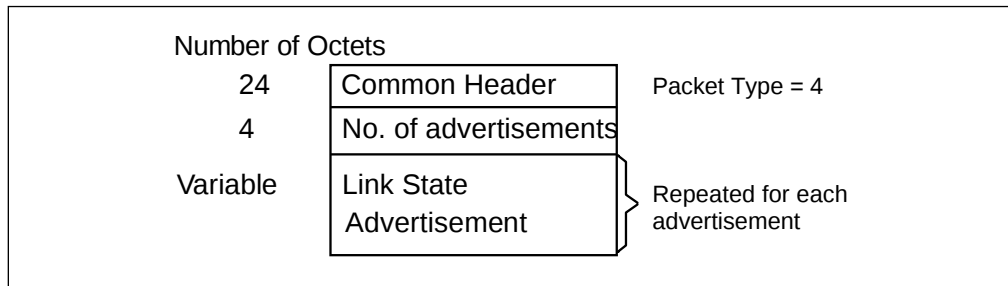


Figure 62. OSPF Link-State Update Packet

Link-state advertisements are of five types: router links, network links, summary links (two types) and AS external links as noted earlier in this section.

Link-state updates pass as a result of link-state requests during database exchange, and also in the normal course of events when routers wish to indicate a change of network topology. Individual link-state update packets can contain multiple link-state advertisements.

It is essential that each OSPF router in an area has the same network topology database, and hence the integrity of link-state information must be maintained.

For that reason link-state update packets must be passed without loss or corruption throughout an area. The process by which this is done is called *flooding*.

A link-state update packet floods one or more link-state advertisements one hop further away from their originator. To make the flooding procedure reliable each link-state advertisement must be acknowledged separately. Multiple acknowledgements can be grouped together into a single *link-state acknowledgement packet*. The format of the link-state acknowledgement packet is shown in Figure 63.




Figure 63. OSPF Link-State Acknowledgement Packet

In order to maintain database integrity it is essential that all link-state advertisements are rigorously checked to ensure validity.

The following checks are applied and the advertisement discarded if:

- The link-state checksum is incorrect
- The link-state type is invalid
- The advertisement's age has reached its maximum
- The advertisement is older than or the same as one already in the database

If an advertisement passes the previous checks, then an acknowledgement is sent back to the originator. If no acknowledgement is received by the originator, then the original link-state update packet is retransmitted after a timer has expired.

Once accepted an advertisement is flooded onward over the router's other interfaces until it has been received by all routers within an area.

Advertisements are identified by their *link-state type*, *link-state ID* and the *advertising router*. They are further qualified by their *link-state sequence number*, *link state age* and *link-state checksum number*.

The age of a link-state advertisement must be calculated to determine if it should be installed into a router's database. Only a more recent advertisement should be accepted and installed. Advertisements are only considered more recent if they have a newer sequence number, if they have the larger checksum (if sequence numbers are equal), or if they have their age set to *max age* (if checksums are equal).

Valid link-state advertisements are installed into the topology database of the router. This causes the topology map or graph to be recalculated and the routing table to be updated.

Link-state advertisements all have a common header. This is shown in Figure 64. The five link-state advertisement types are shown in Figure 65 on page 114, in Figure 66 on page 114, in Figure 67 on page 114, and in Figure 68 on page 115.




Figure 64. OSPF Link-State Header

3.5.2.5 Routing Table Calculation

Each router in an OSPF area builds up a topology database of validated link-state advertisements and uses them to calculate the network map for the area. From this map the router is able to determine the best route for each destination and insert it into its routing table.

Each advertisement contains an age field that is incremented while the advertisement is held in the database. An advertisement's age is never incremented past *max age*. When max age is reached it is excluded from routing table calculation and reflooded through the area as a newly originated advertisement.




Figure 65. OSPF Router Links Advertisement




Figure 66. OSPF Network Links Advertisement




Figure 67. OSPF Summary Links Advertisement




Figure 68. OSPF External Links Advertisement

Routers build up their routing table from the database of link-state advertisements in the following sequence:

1. The shortest path tree is calculated from router and network links advertisements allowing the best routes within the area to be determined.
2. Inter-area routes are added by examination of summary links advertisements.
3. AS external routes are added by examination of AS external links advertisements.

The topology graph or map constructed from the previous process is used to update the routing table. The routing table is recalculated each time a new advertisement is received.

3.5.3 Multicast Extensions to OSPF (MOSPF)

RFC 1584 describes an extension to OSPF Version 2, so that multicast routing capabilities can be added into an OSPF Version 2 routing domain.

IP multicasting is an extension of LAN multicasting to a TCP/IP internet. Multicasting support for TCP/IP hosts has been specified in RFC 1112. In that document, multicast groups are represented by IP class D addresses. Individual TCP/IP hosts join (and leave) multicast groups through the Internet Group Management Protocol (IGMP, also specified in RFC 1112). A host need not be a member of a multicast group in order to send datagrams to the group. Multicast datagrams are to be delivered to each member of the multicast group with the same *best-effort* delivery according to regular (unicast) IP data traffic.

MOSPF provides the ability to forward multicast datagrams from one IP network to another (that is, through Internet routers). MOSPF forwards a multicast datagram on the basis of both the datagram's source and destination (this is sometimes called source/destination routing). The OSPF link-state database provides a complete description of the Autonomous System's topology. By adding a new link-state advertisement (type 6), the *group-membership-LSA*, the location of all multicast group members is pinpointed in the database. The path of a multicast datagram can then be calculated by building a shortest-path tree rooted at the datagram's source. All branches not containing multicast members are pruned from the tree. These pruned shortest-path trees are initially built when the first datagram is received (that is, on demand). The results of the shortest path calculation are then cached for use by subsequent datagrams having the same source and destination.

Routers running MOSPF can be intermixed with non-multicast OSPF routers. Both types of routers can interoperate when forwarding regular (unicast) IP data traffic. Obviously, the forwarding extent of IP multicasts is limited by the number of MOSPF routers present in the Autonomous System (and their interconnection, if any). An ability to *tunnel* multicast datagrams through non-multicast routers is not provided. In MOSPF, just as in the base OSPF protocol, datagrams (multicast or unicast) are routed *as is*; they are not further encapsulated or decapsulated as they transit the Autonomous System.

3.5.3.1 IGMP Interface: The Local Group Database

The local group database keeps track of the group membership of the router's directly attached networks. Each entry in the local group database is a (group, attached network) pair, which indicates that the attached network has one or more IP hosts belonging to the IP multicast destination group. This information is then used by the router when deciding which directly attached networks to forward a received IP multicast datagram onto, in order to complete delivery of the datagram to (local) group members.

The local group database is built through the operation of IGMP. When a MOSPF router becomes a designated router on an attached network, it starts sending periodic IGMP Host Membership Queries on the network. Hosts then respond with IGMP Host Membership Reports, one for each multicast group to which they belong. Upon receiving a Host Membership Report for a multicast group, the router updates its local group database by adding/refreshing the entry. If at a later time reports for a group cease to be heard on the network, the entry is then deleted from the local group database.

It is important to note that on any particular network, the sending of IGMP Host Membership Queries and the listening to IGMP Host Membership Reports is performed solely by the designated router. A MOSPF router ignores Host Membership Reports received on those networks where the router has not been elected the designated router. This means that at most one router performs these IGMP functions on any particular network and ensures that the network appears in the local group database of at most one router. This prevents multicast datagrams from being replicated as they are delivered to local group members. As a result, each router in the Autonomous System has a different local group database.

3.5.3.2 Intra-Area Multicasting

Group-membership-LSAs are specific to a single OSPF area. This means that, just as with OSPF router-LSAs, network-LSAs and summary-link-LSAs, a group-membership-LSA is flooded throughout a single area only. A router attached to multiple areas may end up originating several group-membership-LSAs concerning a single multicast destination, one for each attached area.

Just as in OSPF, each MOSPF area has its own link-state database. The MOSPF database is simply the OSPF link-state database enhanced by the group-membership-LSAs.

3.5.3.3 Inter-Area Multicasting

In OSPF, the area border routers forward routing information and data traffic between areas. In MOSPF, a subset of the area border routers, called the *inter-area multicast forwarders*, forward group membership information and multicast datagrams between areas. Whether a given OSPF area border router is also an MOSPF inter-area multicast forwarder is configuration-dependent.

In order to convey group membership information between areas, inter-area multicast forwarders *summarize* their attached areas group membership to the backbone. This is a very similar functionality to the summary-link-LSAs that are generated in the base OSPF protocol. An inter-area multicast forwarder calculates which groups have members in its attached non-backbone areas. Then for each of these groups, the inter-area multicast forwarder injects a group-membership-LSA into the backbone area.

However, unlike the summarization of unicast destinations in the base OSPF protocol, the summarization of group membership in MOSPF is asymmetric. While a non-backbone area's group membership is summarized to the backbone, this information is not then readvertised into other non-backbone areas. Nor is the backbone's group membership summarized for the non-backbone areas.

To accomplish intra-area multicasting, the notion of *wild-card multicast receivers* is introduced. A wild-card multicast receiver is a router to which all multicast traffic, regardless of multicast destination, should be forwarded. A router's wild-card multicast reception status is per-area. In non-backbone areas, all inter-area multicast forwarders are wild-card multicast receivers. This ensures that all multicast traffic originating in a non-backbone area will be forwarded to its inter-area multicast forwarders and, hence, to the backbone area. Since the backbone has complete knowledge of all areas' group membership, the datagram can then be forwarded to all group members.

3.5.3.4 Interaction with the IGMP Protocol

MOSPF uses the IGMP protocol to monitor multicast group membership. In short, the designated router on a network periodically sends IGMP Host Membership Queries, which, in turn, elicit IGMP Host Membership Reports from the network's multicast group members. These Host Membership Reports are then recorded in the designated router's and backup designated router's local group databases.

Only the network's designated router sends Host Membership Queries. This minimizes the amount of group membership information on the network, both in terms of queries and responses.

Received Host Membership Reports are processed by both the network's designated router and backup designated router. It is the designated router's responsibility to distribute the network's group membership information throughout the routing domain. The backup designated router processes reports so that it too has a complete picture of the network's group membership, enabling a quick cutover upon designated router failure.

3.5.4 3746 IP and OSPF

The 3746 IP router supports the following features of OSPF that have implementation-specific behavior:

- Area border (AB) router support
- Support for stub areas
- Autonomous System border (ASB) support
- OSPF interface support
- Equal-cost multipath routing
- Simple authentication
- OSPF routing policies
- Multicast OSPF (MOSPF) support

There is no support for type of service (TOS) based routing; that is, TOS 0 is the only supported TOS.

For details on configuring 3746 IP OSPF functions, see 4.2.5, “OSPF Configuration on the IBM 3746-9x0” on page 136.

3.5.4.1 Area Border (AB) Router Support

3746 IP supports attachment to multiple areas and summarization of routing information between areas. Area border routers must attach to the backbone (0.0.0.0) and at least one other area. Summarization information from one area will manifest itself as type 3 and 4 link-state advertisements (LSAs) in other areas.

3.5.4.2 Support for Stub Areas

3746 IP supports attachment to stub areas. OSPF Autonomous System external (ASE) LSAs will not be advertised into stub areas. Rather a type 3 network summary LSA for the default route (destination/mask 0.0.0.0) is generated.

3.5.4.3 Autonomous System Border (ASB) Support

3746 IP can be configured as an Autonomous System border (ASB) router. Non-OSPF routes can be imported into OSPF as OSPF Autonomous System externals (ASEs). This implies that the 3746 IP will generate type 5 LSAs.

3.5.4.4 OSPF Interface Support

3746 IP supports the following types of OSPF interfaces:

- Numbered point-to-point (PtP)
Numbered PtP connections are links to which an IP address has been assigned. Examples are ESCON and PPP connections.
- Unnumbered point-to-point (PtP)
Unnumbered PtP connections are links to which no IP address has been assigned. OSPF packets will be sent using the 3746 IP router-ID as source address.
- Broadcast
An example is a token-ring interface. Link-level multicast is used to broadcast OSPF frames to all attached OSPF routers.

- Non-broadcast multi-access (NBMA)

Viewing a network, for example, a frame relay network, as an NBMA network can be used when the network is fully meshed (meaning virtual circuits exist between any pair of routers). In routers that are eligible to become designated routers, neighbors must be configured as well whether or not the neighbor is eligible to become a designated router. The configurable poll interval defines the interval at which the designated router will attempt to contact the neighboring routers to establish an adjacency.

NBMA connections are supported for frame relay networks. If the network is partially meshed, it is more useful to view the network as a point-to-multipoint network.

- Point-to-multipoint (PtM)

This interface type is used to allow NBMA topologies to be non-fully meshed. Rather than electing a designated router for the network and having that router generate network LSAs for the network, each router includes its neighbors in its router LSAs. When the route table is calculated, the network topology will appear as multiple point-to-point links rather than a single cloud. On one side of the frame relay virtual circuit, the neighboring router must be configured to allow the two routers to form an OSPF adjacency.

Point-to-multipoint connections are supported for frame relay networks.

- Virtual link

Virtual links are supported to extend the backbone area's connectivity through a transit area. The two endpoints of the virtual link are area border routers (see 3.5.1.2, "The OSPF Backbone" on page 103).

3.5.4.5 Equal-Cost Multipath Routing

3746 IP supports up to four equal-cost next hops for a route. When multiple next hops exist, the traffic to the destination is spread over the next hops round-robin.

3.5.4.6 Simple Password Authentication

3746 IP supports both simple password and no authentication. When simple password authentication is used, an 8-byte password is included in each OSPF packet. Upon reception, this password is validated, with packets failing validation being dropped.

The authentication type (simple or none) is configured on the area level, while the authentication key is configured for each interface in areas with simple password authentication enabled.

3.5.4.7 OSPF Routing Policies

The 3746 IP OSPF policy can be explained in terms of rules for:

- Advertisements of OSPF routes
- Import of external routes into OSPF as OSPF AS external (ASE) routes
- Generation of the default route and import as an OSPF ASE LSA
- OSPF route policy when multiple routes to a destination exist

OSPF Advertisement of Routes: OSPF allows filtering of LSAs on area boundaries only. All routers within an area have the same view of the area topology.

To limit the number of LSAs advertised outside the area, one can configure the network ranges associated with an area at the area boundary. This, in effect, will aggregate a number of networks into a single advertisement. The cost associated with the network range will be the lowest for any of the component networks. Additionally, one can define a network range that will not be visible.

Normally, OSPF ASE LSAs are flooded throughout the entire routing domain. One can prevent this for a given area by defining the area as a stub area. Within the stub area, the area border router will advertise a single default route (destination/mask 0.0.0.0).

Importing Routes into OSPF: 3746 IP allows configuration whether or not non-OSPF routes should be imported as OSPF AS external routes and advertised as OSPF ASE LSAs throughout the OSPF routing domain.

Imported routes can be imported as ASE type 1 (router) or ASE type 2 (network) routes. ASE type 1 routes always override type 2 routes. It has a single metric, which is the sum of the path cost to the AS border router, and the AS border router metric. Conversely, the metric for a type 2 ASE has an internal and external components: the path cost to the AS border router and the route's external metric. When comparing ASE type 2 routes to the same destination, the one with the lower external metric will always be preferred, independent of the internal metric. The metric used for both the OSPF ASE type 1 route and the external component of the type 2 route is the metric from the protocol from which it is imported.

Default Route: 3746 IP allows generation of default routes and advertisement of the routes through an AS external (ASE) LSA. For details, see 4.5, "Routing Protocols Interoperability" on page 146.

3.6 Border Gateway Protocol (BGP)

The *Border Gateway Protocol* (BGP) Version 4 is an exterior gateway protocol defined in RFC 1654 (made obsolete by RFC1771). A companion document, RFC 1656 - *BGP-4 Protocol Document Roadmap and Implementation Experience* (made obsolete by RFC1773), details experience of its use in the Internet.

BGP was built upon experience with EGP and is intended as a functionally superior replacement for EGP. BGP Version 4 is an inter-Autonomous System gateway protocol. EGP has no ability to determine routes within an AS, it must rely on an interior gateway protocol, such as RIP or OSPF, to do this. EGP is an IAB standard protocol; its status is recommended.

BGP was introduced in the Internet for the loop-free exchange of routing information between Autonomous Systems. Based on Classless Inter-Domain Routing (CIDR), BGP has since evolved to support the aggregation and reduction of routing information.

In essence, CIDR is a strategy designed to address the following problems:

- Exhaustion of class B address space
- Routing table growth

CIDR eliminates the concept of address classes and provides a method for summarizing multiple different routes into single routes. This significantly reduces the amount of routing information that BGP routers must store and exchange.

3.6.1 Introduction

BGP is not a routing protocol, but a reachability protocol. In essence, BGP routers selectively collect and advertise reachability information to and from BGP neighbors in their own and other Autonomous Systems. Reachability information consists of the sequences of AS numbers that form the paths to particular BGP speakers and the list of IP addresses that can be reached via each advertised path. An AS is an administrative group of networks and routers that share reachability information using one or more interior gateway protocols (IGPs), such as RIP or OSPF.

Routers that run BGP are called BGP speakers. These routers function as servers with respect to its BGP neighbors (its clients). Each BGP router opens a connection and listens for incoming connections from neighbors at this well-known address. The router also opens active TCP connections to enabled BGP neighbors. This TCP connection enables BGP routers to share and update reachability information with neighbors in the same or other Autonomous Systems.

Connections between BGP speakers in the same AS are called *internal* connections, while connections between BGP speakers in different Autonomous Systems are *external* connections. A single AS may have one or many BGP connections to outside Autonomous Systems.

Figure 69 shows three Autonomous Systems. The BGP speakers in AS-1 establish an internal connection with each other and an external connection with their neighbor in AS-2 and AS-3, respectively. Once the connections have been established, the routers will be able to share reachability information.




Figure 69. BGP Connections

3.6.1.1 Originate, Send and Receive Policies

Decisions on which reachability information to advertise (send) and which to accept (receive) are made on the basis of explicitly defined policy statements. IBM's BGP implementation supports three types of policy statements:

- Originate Policies

The originate policies decide which interior gateway protocol routes are advertised to BGP neighbors.

- Send Policies

The send policies decide which BGP received routes are forwarded to BGP neighbors.

- Receive Policies

The receive policies decide which BGP received routes are injected into the interior gateway protocol routing tables.




Figure 70. Route Aggregation and AS Policies

Once a TCP connection is established, the BGP speaker in AS-1, shown in Figure 70, can send its entire routing table to its BGP neighbor in AS-2. However, for security or other reasons, it may not be desirable to send reachability information on each network to AS-2. Similarly, it may not be desirable for AS-2 to receive reachability information on each network in AS-1.

In addition to deciding which routing information is sent, the BGP speaker also has the ability to condense its routing information. As an example, assume that AS-1 contains multiple class C networks. Instead of sending separate routes for each individual network, an *aggregated* route comprising all networks can be sent. For details, see 3.6.2.4, “Route Aggregation” on page 127.

3.6.2 Protocol Description

The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems. This network reachability information includes information on the list of Autonomous Systems (ASs) that reachability information traverses. This information is sufficient to construct a graph of AS connectivity from which routing loops may be pruned and some policy decisions at the AS level may be enforced.

BGP-4 provides a new set of mechanisms for supporting Classless Inter-Domain Routing. These mechanisms include support for advertising an IP prefix and eliminates the concept of network *class* within BGP. BGP-4 also introduces mechanisms that allow aggregation of routes, including aggregation of AS paths. These changes provide support for supernetting schemes.

BGP runs over a reliable transport protocol. This eliminates the need to implement explicit update fragmentation, retransmission, acknowledgement, and sequencing. Any authentication scheme used by the transport protocol may be used in addition to BGP's own authentication mechanisms. The error notification mechanism used in BGP assumes that all outstanding data will be delivered before the connection is closed.

BGP uses TCP as its transport protocol. TCP meets BGP's transport requirements and is present in virtually all commercial routers and hosts. BGP uses TCP port 179 for establishing its connections.

BGP systems form a transport protocol connection between one another. They exchange messages to open and confirm the connection parameters. The initial data flow is the entire BGP routing table. Incremental updates are sent as the routing tables change. BGP does not require periodic refresh of the entire BGP routing table. Therefore, a BGP speaker must retain the current version of the entire BGP routing tables of all of its peers for the duration of the connection. Keep alive messages are sent periodically to ensure the liveliness of the connection. Notification messages are sent in response to errors or special conditions. If a connection encounters an error condition, a notification message is sent and the connection is closed.

If a particular AS has multiple BGP speakers and is providing transit service for other ASs, then care must be taken to maintain a consistent view of routing within the AS. A consistent view of the interior routes of the AS is provided by the interior routing protocol. A consistent view of the routes exterior to the AS can be provided by having all BGP speakers within the AS maintain direct BGP connections with each other. Using a common set of policies, the BGP speakers arrive at an agreement as to which border routers will serve as exit/entry points for particular networks outside the AS. This information is communicated to the AS's internal routers, possibly via the interior routing protocol. Care must be taken to ensure that the interior routers have all been updated with transit information before the BGP speakers announce to other ASs that transit service is being provided.

Connections between BGP speakers of different ASs are referred to as *external* links. BGP connections between BGP speakers within the same AS are referred to as *internal* links. Similarly, a peer in a different AS is referred to as an external peer, while a peer in the same AS may be described as an internal peer.

The BGP protocol comprises four main stages:

- Opening and confirming a BGP connection with a neighbor router
- Maintaining the BGP connection
- Sending reachability information
- Notification of error conditions

3.6.2.1 Opening and Confirming a BGP Connection

BGP communications between two routers commences with the TCP transport protocol connection being established. A description of this is outside the scope of this document, but can be found in *TCP/IP Tutorial and Technical Overview*, GG24-3376.

Once the connection is established, each router sends an *open* message to its neighbor.

The BGP open message, like all BGP messages, consists of a standard header plus packet type specific contents. The standard header consists of a 16-byte *marker* field, which is set to all ones, the *length* of the total BGP packet, and a *type* field that specifies the packet to be one of four possible types:

1. Open
2. Update
3. Notification

4. Keep alive

The format of the BGP header is shown in Figure 71.




Figure 71. BGP Message Header

The open message defines the originating router's *AS number*, its BGP *router identifier* and the *hold time* for the connection. If no keep alive, update or notification messages are received for a period of hold time, the originating router assumes an error, sends a notification message, and closes the connection.

The open message also provides an *authentication code* and *authentication data*. The use of these fields is not fully defined in the RFC and current BGP implementations use authentication code 0 with authentication data of all zeros.

The format of the open message is shown in Figure 72.




Figure 72. BGP Open Message

An acceptable open message is acknowledged by a *keep alive* message. Once neighbor routers have sent keep alives in response to opens, they can proceed to exchange further keep alives, notifications and updates.

3.6.2.2 Maintaining the BGP Connection

BGP messages must be exchanged periodically between neighbors. If no messages are received for a period defined by hold time in the open message, then an error on the connection is assumed.

BGP uses keep alive messages to maintain the connection between neighbors. Keep alive messages consist of the BGP packet header only with no data. The RFC recommends that they should be sent at intervals of approximately one third of hold time.

3.6.2.3 Sending Reachability Information

Reachability information is exchanged between BGP neighbors in update messages. Update messages are used to transfer routing information between BGP peers. The information in the update packet can be used to construct a graph describing the relationships of the various Autonomous Systems. By applying rules to be discussed, routing information loops and some other anomalies may be detected and removed from inter-AS routing.

An update message is used to advertise a single feasible route to a peer, or to withdraw multiple infeasible routes from service. An update message may simultaneously advertise a feasible route and withdraw multiple unfeasible routes from service. The update message always includes the fixed-size BGP header and can optionally include the other fields as shown in Figure 73.





Figure 73. BGP Update Message

The *unfeasible routes length* is a 2-octet field that indicates the total length of the withdrawn routes field. A value of 0 indicates that no routes are being withdrawn from service, and that the withdrawn routes field is not present in this update message.

The *withdrawn routes* field is a variable length field that contains a list of IP address prefixes for the routes that are being withdrawn from service. Each IP address prefix is encoded as a (length, prefix) pair.



The length field indicates the length in bits of the IP address prefix. A length of zero indicates a prefix that matches all IP addresses. When the length field is zero, no prefix field is included.

The prefix field contains IP address prefixes followed by enough trailing bits to make the end of the field fall on an octet boundary. Note that the value of trailing bits is irrelevant.

The *total path attribute length* is a 2-octet unsigned integer that indicates the total length of the path attributes field in octets. A value of 0 indicates that no network layer reachability information field is present in this update message.

Each *path attribute* consists of a triple set of values: *attribute flags*, *attribute type* and *attribute value*. Three of the attribute flags provide information about the status of the attributes types and may be *optional* or *well-known*, *transitive* or *non-transitive* and *partial* or *complete*.

Attribute flags must be read in conjunction with their associated attribute types. There are five attribute types which together define an advertised route.

The first three of these attributes are mandatory and must be supplied with all network advertisements. Each attribute type sent in a BGP update message must have its attribute flags set as indicated, otherwise an error occurs.

The format of BGP path attributes is shown in Figure 74.




Figure 74. BGP Path Attributes

The remaining octets of the path attribute represent the attribute value and are interpreted according to the attribute flags and the attribute type code. The supported attribute type codes, their attribute values and uses are the following:

- Origin (Type Code 1)

A well-known mandatory attribute that defines the origin of the route as an interior gateway protocol, an exterior gateway protocol or other (for example a static route).

- AS_Path (Type Code 2)

AS_Path is a well-known mandatory attribute that is composed of a sequence of AS path segments. Each AS path segment is represented by a triple segment (path segment type, path segment length, and path segment value).

The *path segment* type is a 1-octet long field with the following values defined:

- 1 - AS_SET, an unordered set of ASs a route in the update message has traversed
- 2 - AS_SEQUENCE, an ordered set of ASs a route in the update message has traversed

The *path segment* length is a 1-octet long field containing the number of ASs in the path segment value field.

The *path segment* value field contains one or more AS numbers, each encoded as a 2-octet long field.

- Next_Hop (Type Code 3)

This field contains the IP address of the border router that should be used as the next hop to the destinations listed in the Network Layer Reachability field of the update message.

- Multi_Exit_Disc (Type Code 4)

This is an optional nontransitive attribute that is a 4-octet non-negative integer. The value of this attribute is used to discriminate among multiple exit points to a neighboring Autonomous System.

- Local_Pref (Type Code 5)

Local_Pref is a well-known discretionary attribute that is a 4-octet non-negative integer. It is used by a BGP speaker to inform other BGP speakers in its own Autonomous System of the originating speaker's degree of preference for an advertised route.


- Atomic_aggregate (Type Code 6)

Atomic_aggregate is a well-known discretionary attribute of length 0. It is used by a BGP speaker to inform other BGP speakers that the local system selected a less specific route without selecting a more specific route that is included in it.

- Aggregator (Type Code 7)

Aggregator is an optional transitive attribute of length 6. The attribute contains the last AS number that formed the aggregate route (encoded as 2 octets), followed by the IP address of the BGP speaker that formed the aggregate route (encoded as 4 octets).

The *Network Layer Reachability Information* is a variable length field containing a list of IP address prefixes. Each IP address prefix is encoded as a (length, prefix) pair.



The Prefix field contains IP address prefixes and is followed by enough trailing bits to make the end of the field fall on an octet boundary.

An update message can advertise at most one route, which may be described by several path attributes. All path attributes contained in a given update message apply to the destinations carried in the Network Layer Reachability Information field of the update message.

3.6.2.4 Route Aggregation

Previous versions of BGP did not support supernetting nor subnetting as update messages being sent contained 32-bit Internet addresses that represented the (class A, B, or C) network that was reachable. That is, older BGP versions are *class-oriented*.

In BPG Version 4, however, networks are represented by a two-component structure (see, for example, the Network Layer Reachability Information field described in the previous section).

Number of Octets	
1	Length
Variable	Prefix

The Length field indicates the number of bits in the address prefix and is followed by the number of octets to hold the prefix.

When BGP speakers are exchanging routing information, they not only exchange routes but they also condense (that is *aggregate*) routes. Route aggregation leads to smaller routing tables.




Figure 75. Route Aggregation

Figure 75 depicts an Internet consisting of five ASs. AS-1 and AS-2 each use a single class C subnet, 197.8.2. and 197.8.3, respectively. The routes from AS-3 to AS-1 and AS-2 can be represented as a 24-bit prefix of 197.8.2 and 197.8.3, respectively.

AS-3 uses two class C subnets, 197.8.0. and 197.8.1. The routes from AS-1 and AS-2 to AS-3 can be represented as a 23-bit prefix of 197.8.0. Hereby, two class C subnet routes are aggregated in a single route.

Maximum route aggregation can be accomplished when AS-3 informs AS4 about the routes to AS-3 itself, AS-1, and AS-2. Instead of announcing three paths, AS-3 forwards a route consisting of a 22-bit prefix of 197.8.0. In addition to the network information, path attributes are sent to indicate the numbers of the ASs that are reachable.

Note: As can be seen in this example, route aggregation strongly relies on allocating consecutive address ranges within ASs.

3.6.2.5 Notifying Errors

Notification messages are sent to a neighbor router when error conditions are detected. The BGP transport connection is closed immediately after a notification message has been sent.

Notification messages consist of an *error code* and an *error subcode* that further qualifies the main error. The format of notification messages is shown in Figure 76 on page 129.




Figure 76. BGP Notification Message

Error codes that are provided by BGP are as follows:

- Message Header Error
- Open Message Error
- Update Message Error
- Hold Timer Expired
- Finite State Machine Error
- Cease

A *data field* is included in the notification message to provide additional diagnostic information.

3.6.3 3746 IP and BGP Version 4

The IBM 3746-9x0 supports a full implementation of BGP Version 4 with *null authentication*. Only BPG Version 4 support is provided, and earlier BPG versions are not supported.

Configuration options exist to:

- Enable BGP and specify the local Autonomous System number
- Define BGP neighbors
- Define (Exclude) ASs from which no routing information will be accepted
- Define send, receive, and originate policies

See 3.6.1.1, "Originate, Send and Receive Policies" on page 121.

- Define aggregate routes

The 3746 IP router requires that aggregated routes (see Figure 75 on page 128) are preconfigured. When defining aggregated routes make sure that the individual routes that make up the aggregated route are not exported (that is, define a send policy).

For details on configuring 3746 IP BGP functions, see 4.3.2, "BGP Configuration on the IBM 3746-9x0" on page 141.

Chapter 4. IBM 3746-9x0 and IP Routing Protocols

This chapter provides hints and tips to designers of IP internetworks to allow them to make the best use of the IBM 3746-9x0.

It provides general information about the limitations of each protocol and when they should and should not be used. It also provides specific information to assist network designers in configuring each of the IP routing protocols on the IBM 3746-9x0.

Many networks demand the use of multiple interior gateway protocols or the use of interior and exterior gateway protocols in combination. When multiple protocols are used it is necessary to *export* routes from one protocol to another. The IBM 3746-9x0 provides export capability and this is also covered in this chapter.

Chapter 7, "Using CCM - IP Configuration Examples" on page 283 details configuration examples for the 3746 IP router. Many references are made to this chapter to show the practical implementation of design considerations and configuration guidelines discussed.

4.1 Using RIP

RIP is currently the most common IP routing protocol. This is because it has been available with the UNIX operating system for some time and, because it is straightforward to implement, it has been ported to or re-implemented on other systems. It is supported by all major router vendors.

Most RIP networks are relatively small, and many have evolved to meet purely local requirements. In these networks general purpose computer systems are often used as convenience routers, and RIP has been used because such systems generally only support that single routing protocol.

There are major limitations to RIP. These limitations, expectedly, are highlighted when large internetworks are to be designed. They were discussed in 3.3, "Routing Algorithms" on page 92 but are listed below for completeness:

- No support for variable length subnet masks
- Limit of 15 hops in a network
- No support for alternate routes based on IP type of service
- Slow convergence on large internetworks (unless triggered updates supported, as is the case in the IBM 3746)
- Routing tables broadcast even if unchanged
- Routing tables can be very large as they contain all routes
- No support for discontinuous networks

Note: A discontinuous network is a (class A, B, or C) network split up into multiple parts by one or more other networks.

RIP is a very robust protocol and there is a very high probability that implementations from different vendors will interwork without difficulty, but due to

the previous limitations, it is not recommended that you use RIP as the major routing protocol for new network designs.

RIP should only be used in situations where it is necessary to interwork with an existing network already using RIP. If an existing internetwork is growing in size it is recommended, where possible, that the extensions be implemented using OSPF. This can be done by migrating the entire network to OSPF or by running parts of the network with OSPF and parts with RIP. If this is not possible an OSPF backbone could be implemented to link together the existing RIP-based networks.

If RIP is used in a network, the designer should be aware of its limitations and attempt to configure routers to reduce their impact to the minimum.

In order to optimize the use of RIP there are several guidelines that should be adopted when interface settings are configured. Their effect will be to minimize unnecessary network traffic and minimize RIP processing overhead.

- Do not enable RIP unless there is at least one interface on which it is used.
- Enable an IP interface for RIP only when there are RIP partners reachable via that interface.
- If possible, do not enable RIP for interfaces to low-speed point-to-point networks or networks (for example, frame relay) for which usage charges apply. Instead use static or default routes to define partner routers.

4.1.1 RIP Configuration on the IBM 3746-9x0

On the IBM 3746-9x0, RIP is, by default, disabled and needs to be enabled before it can be used (see Figure 77). Configuration options exist to control the broadcast of a Default Route.

ITSO/3746-9x0/RIP - General Parameters

☒ Enable RIP

Originate Default Route

☐ Always originate default route

☐ Originate default route if BGP routes available

From AS number: numerical [1-65535]

To network number:

☐ Originate default route if OSPF routes available

Default route cost: numerical [1-16]

Route Acceptance

Network address:

Figure 77. RIP - General Parameters

RIP will always accept the net, subnet, and host routes defined within Route Acceptance. If no Route Acceptance routing information has been entered, the receipt of RIP routes is controlled by the parameters set in RIP - Parameters Per IP Address (see Figure 78 on page 133).

After enabling RIP, it needs to be enabled on all IP interfaces on which RIP functions are required. For 3746 IP routers that are part of multiple Autonomous Systems (ASs), interfaces must be assigned to a specific AS. Configuration options exist per interface to accept net, subnet, and/or host routes and to override static and/or default routes. When the receipt of net, subnet, and/or host routes has been disabled, RIP will only accept routes defined within Route Acceptance in the RIP - General Parameters panel (see Figure 77 on page 132).

ITS0/3746-9x0/OSPF/RIP - Parameters Per IP Address

IP Address 192.168.202.8

☐ Add OSPF (O) ☒ Add RIP (R)

OSPF parameters... RIP parameters... OSPF neighbors (N)...

Select An IP Address

DLC	Port	Port name	Dial circuit	IP address	O	R	N
TR	2080	PORT2080		CBSP automatic			
TR	2144	PORT2144		192.168.150.3			
TR	2144	PORT2144		9.24.104.168			
TR	2176	PORT2176		192.168.157.2		X	
TR	2208	PORT2208		10.0.0.3			
ESCON	2240	HL2240IP		192.168.202.8		X	

OK Search... Sort... Cancel Help

Figure 78. RIP - Parameters Per IP Address

Once an interface is enabled for RIP, the 3746 IP router will broadcast RIP packets. Configuration options exist to send net, subnet, routes, static, and/or default routes. The use of poison reverse (see 3.3.2.2, "Split Horizon with Poison Reverse" on page 97) is, by default, activated. Be aware that poison reverse on the 3746-9x0 is mutually exclusive with the use of split horizon.

Note: The changing of this default is not supported by CCM, it must be done via telnet. To disable sending poison reverse routes, connect to the NNP via telnet, enter CONFIG, select *Protocol IP*, then enter the following command:

```
DISABLE SENDING POISONED_REVERSE_ROUTES
```

or,

```
ENABLE SENDING POISONED_REVERSE_ROUTES
```

You will be prompted for the address of the IP interface on which you wish to disable or enable poison reverse. Disabling poison reverse will automatically enable split horizon on the interface.

4.2 Using OSPF

OSPF is a relatively new routing protocol and, for that reason, does not have the installed base of RIP. However, most large internetworks currently in design or being implemented use OSPF as the major routing protocol. This is because OSPF overcomes many of the limitations of RIP:

- It supports variable length subnet masks.
- It supports alternate routes based on IP type of service (not supported by the IBM 3746 Network Processor).
- It supports equal-cost multipath routes.
- There is no hop count limit.
- Network convergence is fast and can be designed for by use of small areas.
- Only network topology changes are advertised thus reducing network traffic.

Our recommendation is to use OSPF as the major routing protocol for all new network designs and replace RIP routed networks when possible. If this cannot be done, an OSPF backbone should be used to connect the existing RIP routed networks.

When designing an OSPF network, consideration should be given to partitioning the network into backbone and non-backbone areas. Selecting optimum area sizes can significantly reduce the amount of OSPF traffic passing over individual links. This can be important if links are low-speed point-to-point.

Areas may also be defined as stub areas such that OSPF ASE link advertisements are not flooded into them. If areas have single exit points and the AS connects to public internets, then the use of stub areas can be of great benefit.

Consideration should also be given to the IP address ranges being used within each area. The 3746-9x0 allows users to define the address range being used within an area. Using consecutive subnets allows route aggregation (supernetting) to limit the number of routes advertised and reduces the size of routing tables in routers in other areas. For details, see "OSPF Advertisement of Routes" on page 119.

4.2.1 OSPF Areas

One of the key design decisions to be made for an OSPF network is how it should be partitioned into *areas*. All OSPF networks must consist of at least one area, the backbone, plus other areas that use the backbone for inter-area traffic.

Many small to medium-sized networks may consist of just a backbone area. In this case all routers are configured as part of area 0.0.0.0 and there will be no *area border routers*. If there are default routes or if exterior gateway protocols are being used to communicate with other ASs, then one or more routers must act as AS boundary routers.

Backbone-only networks are acceptable only up to a certain size. As area size increases the number of link-state advertisements that must be flooded over all links increases and, particularly when point-to-point links are concerned, there is a time when the network must be partitioned.

It is difficult to define when a network reaches the size where it must be configured into multiple areas. It is dependent on overall network topology, speed of links and administrative factors.

Some possibly useful guidelines that can be used in making this decision are as follows:

- OSPF has a design point of 100 routers per area maximum.
- OSPF areas should be defined, where possible, in line with existing administrative boundaries.
- OSPF areas should be defined such that local administration errors will have a minimum impact on the overall network.
- OSPF areas should be defined to minimize traffic over point-to-point links.
- OSPF areas should, where possible, be defined to keep the backbone contiguous and hence avoid virtual links.

Once the decision is made on area partitioning, then routers can be configured with area information.

Each area must be given an *area identifier*. The backbone must use identifier 0.0.0.0, and other areas conventionally use an identifier corresponding to the network number for the networks or subnetworks included in this area.

In addition, areas must be defined as *transit* or *stub*. The normal choice for an area would be transit (the backbone *must* be transit), although stub can be chosen for areas with single area border router exit points. The advantage of configuring an area as a stub is that a default route is used to define the exit point from the stub and no AS external link advertisements are flooded into it. This can significantly reduce OSPF link-state information within an area if large numbers of external routes are available to the OSPF network.

The final choice that must be made for an area is the *authentication scheme* to be used for all OSPF message exchanges within the area. The OSPF common header includes two fields, *authentication type* and *authentication data*. Two authentication types are possible: *none*, in which case no authentication data is passed in the data field, or *simple*, in which case an eight-byte password is passed in the data field. All routers within an area must be configured with the same authentication type. All routers on a network must use the same password value.

4.2.2 OSPF Interfaces

OSPF area configuration for a router requires that each interface to an area be defined by its IP address and by a number of other parameters. Care should be taken with area border routers to ensure that interfaces are assigned to the correct area.

All router interfaces to the same network must be assigned the same passwords if an authentication type of *simple* is chosen. If not, OSPF messages will be rejected. All router interfaces to the same network must be assigned the same timer values

for *HelloInterval* and *RouterDeadInterval*. The *HelloInterval* is the length of time in seconds between Hello packets on the interface, and the *RouterDeadInterval* is the number of seconds before a neighbor declares the router down if no Hellos are received. These two parameters are sent in all Hello messages. If they are not configured with the same value as in neighbor routers then the Hello messages will be rejected. It is essential, therefore, that all routers on a network use the same values. It is recommended, for this reason, that the defaults be taken wherever possible.

Each router interface must be assigned a *cost* based on the type of network link. This cost is used in the determination of the total path cost to a destination.

If the interface is to a non-broadcast network it is necessary to configure the IP addresses of neighbor routers on the network. This is to allow the router to send OSPF messages to neighbors using specific IP addressing.

4.2.3 Area Ranges

Area ranges need only be configured on area border routers. They represent the address ranges that are contained within the area and which will be summarized in summary links advertisements to other areas attached to the area border router.

Configuration of area ranges is optional. When using (consecutive) address ranges that can be aggregated in a single route, or a limited number of routes, defining area ranges helps to minimize the number of routes advertised, thereby reducing the routing table within routers in other areas.

4.2.4 Virtual Links

Virtual links are necessary to connect non-contiguous sections of an OSPF backbone. Virtual links are configured between two area border routers through a non-backbone area by specifying the IP address of the destination area border router (the virtual neighbor identifier) and the area through which the virtual link will pass. They are configured for the backbone area 0.0.0.0 only.

Note: Careful design of network topology should be used to avoid the use of virtual links where possible.

4.2.5 OSPF Configuration on the IBM 3746-9x0

On the IBM 3746-9x0, OSPF is, by default, disabled and needs to be enabled before it can be used (see Figure 79 on page 137). After enabling OSPF, the areas to which the router attaches have to be defined, and OSPF needs to be enabled on any of the IP interfaces on which OSPF functions are required. When using the OSPF multicast facility (see 3.5.3, "Multicast Extensions to OSPF (MOSPF)" on page 115), the multicast group address can be defined. Make sure these are valid class D addresses.

ITS0/3746-9x0/OSPF - General/Multicast Parameters

General OSPF Parameters

☒ Enable OSPF

Number of AS external routes: 100 numerical [1-4000]

Number of OSPF routers: 50 numerical [1-4000]

Multicast

☐ Enable intra-area multicasting

☐ Enable inter-area multicasting

Group Addresses

Group address:

Add

Modify

Delete

OK Cancel Help

Figure 79. OSPF - General/Multicast Parameters

All OSPF areas to which the 3746-9x0 connects need to be defined (see Figure 80). The area number must be unique and consistent with the area numbers defined on other routers within the area. The area number for the OSPF backbone is 0.0.0.0. All routers within an area must define the same values for Stub Area and Authentication. On stub area border routers, the import of Import (route) summaries can be disabled, if proper default routes have been defined.

ITS0/3746-9x0/OSPF - Area Configuration

Configure an OSPF Area

Area number: 0.0.0.0

☐ Stub area ☐ Authentication (A)

Stub Area

Default cost: 1 num [1-65535]

☐ Import summaries (IS)

OSPF Areas Already Configured

Area number	A	Cost	Stub IS
0.0.0.0			

Delete

Ranges...

OK Cancel Help

Figure 80. OSPF - Area Configuration

OSPF area border routers may define address ranges (see Figure 81 on page 138). The address ranges entered may be a list of subnet address ranges within the same area or, for example, a single route in which all subnet address ranges have been aggregated. The latter results in fewer route advertisement and reduces the routing tables in routers within other areas. Hybrid solutions, specifying both single subnet and aggregated routes, are possible as well.

The screenshot shows a window titled "ITSO/3746-9x0/OSPF - Area Ranges". Inside the window, there are two main sections:

- Configure an Area Range**: This section contains fields for "IP address:" and "Subnet mask:", each followed by a dotted grid input area. Below these is a checkbox labeled "Advertise (A)" which is checked. To the right of this section are three buttons: "Add", "Modify", and "Delete".
- Area Ranges Already Configured**: This section features a table with three columns: "IP address", "Subnet mask", and "A". The table body consists of a large dotted grid. A vertical scrollbar is located on the right side of the grid.

At the bottom of the window are three buttons: "OK", "Cancel", and "Help". An arrow points to the "OK" button.

Figure 81. OSPF - Area Ranges

Once an interface is enabled for OSPF, the OSPF parameters per IP address (see Figure 82 on page 139) must be set. Make sure that the proper area number is defined. The authentication key, the Hello and Dead router intervals must be the same within any area. The Cost value allows you to prioritize interfaces. The Priority field plays a role during the election of the designated router. For details on the interface parameters, see 4.2.2, “OSPF Interfaces” on page 135. To simplify the configuration process, use defaults when possible.

ITS0/3746-9x0/OSPF/RIP - Parameters Per IP Address

IP Address 192.168.157.2

☒ Add OSPF [O] ☐ Add RIP [R]

OSPF parameters... RIP parameters...

OSPF neighbors (N)...

Select An IP Address

DLC	Port	Port name	Dial circuit	IP address	O	R	N
TR	2080	PORT2080		CBSP automatic			
TR	2144	PORT2144		192.168.150.3			
TR	2144	PORT2144		9.24.104.168			
TR	2176	PORT2176		192.168.157.2		X	
TR	2208	PORT2208		10.0.0.3			
ESCON	2240	HL2240IP		192.168.202.8			

OK Search... Sort... Cancel Help

Figure 82. OSPF - Parameters Per IP Address

On non-broadcast networks (for example, frame relay) it is necessary to configure the IP addresses of neighbor routers on the network. This is to allow the router to send OSPF messages to neighbors using specific IP addressing. Indicate, hereby, which neighbors are eligible to become the designated router.

In the situation that your backbone area is non-contiguous, virtual links (see Figure 83 on page 140) have to be defined. For details on OSPF virtual links, see 4.2.4, "Virtual Links" on page 136.

Figure 83. OSPF - Virtual Links

4.3 Using BGP Version 4

BGP Version 4 is the newer and the preferred exterior gateway protocol. It was developed on the basis of experience gained with EGP in the Internet. It, therefore, offers several enhancements to the features offered by EGP. In particular it provides:

- Advertisement of all routes used by a BGP router, not only those within the local AS
- Full paths to destinations that can be through multiple ASs
- Information about how a route was learned by a BGP router
- Next hop information, which need not be the neighbor router on the same network
- Intra-area metrics to allow selection to be made between alternate routes to a destination
- Support for neighbor routers that are not on the same common network
- Support for *classless* routes that allows route aggregation leading to smaller routing tables

Provision of this additional information by an originating router allows complex policy decisions to be made based on the information.

BGP-4 should be considered when you wish to use public IP services to communicate with other organizations whose networks reside in separate ASs. This requires that you have registered your organization, have been issued an AS number, and have registered the network numbers used within your organization

Using BGP requires that routers on both ends of a BGP connection support the same BPG version (that is, Version 4).

4.3.1 AS Numbers

BGP can only be configured when the local AS has been allocated an *AS number*, and the identity of the ASs to which neighbor BGP routers belong is known. If this information is not available then it is not possible to proceed. ASs are designated numerically in the range 1 to 65534. AS numbers for an organization must be registered publicly before they can be used.

AS numbers are used to qualify the local router originating BGP packets and also to fully define a neighbor router whose routes are being advertised. AS numbers are further used as filters for routes exported between BGP and interior gateway protocols.

4.3.2 BGP Configuration on the IBM 3746-9x0

On the IBM 3746-9x0, BGP is, by default, disabled and needs to be enabled before it can be used (see Figure 84). After enabling BGP, the (local) AS number must be defined. In addition, (remote) AS numbers can be defined. Routes from these ASs will be excluded from the local routing table.

The screenshot shows a dialog box titled "ITSO/3746-9x0/BGP - General Parameters/Exclu". It is divided into two main sections. The first section, "General Parameters", contains a checked checkbox for "Enable BGP", an unchecked checkbox for "Send subnet routes", a text field for "AS number" with the value "1", and a text field for "TCP segment size" with the value "1024". The second section, "Excluded Autonomous Systems (AS)", contains a text field for "AS no. to exclude" with the value "1". Below this field is a list box with a grid pattern, and to its right are three buttons: "Add", "Modify", and "Delete". At the bottom of the dialog are three buttons: "OK", "Cancel", and "Help".

Figure 84. BGP - General Parameters

After BGP has been enabled, *BGP neighbors* must be defined (see Figure 85 on page 142). BGP neighbors can be of two types: *enabled* and *disabled*. Enabled neighbors are fully defined and communication with them is initiated on BGP startup. Communications with disabled neighbors can be initiated by the neighbor only. BGP neighbors are defined by their IP address. External neighbors must be adjacent (that is, attach to the same (sub)net), internal neighbors do not have to be adjacent.

ITS0/3746-9x0/BGP - Neighbors

Configure a BGP Neighbor

IP address: 10.0.0.8 ☒ Enable neighbor

AS number: 2 numerical [1-65535]

Initialization timer: 12 seconds [0-65535]

Connect retry timer: 120 seconds [0-65535]

Hold timer: 90 seconds [0-65535]

TCP segment size: 1024 bytes [28-65535]

Add Modify

BGP Neighbors Already Configured

IP address	Status	AS number
10.0.0.8	Enable	2

Delete

OK Cancel Help

Figure 85. BGP - Neighbors

Figure 86 shows how the BGP receive policies can be defined. The receive policies control which BGP-received routes will be injected into the router routing table and forwarded into the RIP and/or OSPF domain to which the 3746 IP router connects.

ITS0/3746-9x0/BGP - Receive Policies

Configure a BGP Receive Policy

Policy Type: ☒ Inclusive (I) ☐ Exclusive (E)

Address Match (AM): ☐ Exact (E) ☒ Range (R)

Network IP address: 0.0.0.0 Network mask: 0.0.0.0 IGP metric: 1 num. [1-65535]

Originating AS number: 0 num. [0-65535] Adjacent AS number: 0 num. [0-65535]

Add Modify

BGP Receive Policies Already Configured

No.	Type	AM	Network IP address	Network mask	Originating/Adjacent AS numbers
1	I	R	0.0.0.0	0.0.0.0	0 0

Delete Move up Move down

OK Save as defaults Cancel Help

Figure 86. BGP - Receive Policies

BGP can be configured to permit the import of the default route (0.0.0.0) from BGP neighbors. Before enabling this option, make sure you understand the operational consequences.

Figure 87 shows how the BGP send policies can be defined. The send policies control which BGP-received routes are forwarded to BGP neighbors.

ITS0/3746-9x0/BGP - Send Policies

Configure a BGP Send Policy

Policy Type: ☒ Inclusive [I] ☐ Exclusive [E]

Address Match [AM]: ☐ Exact [E] ☒ Range [R]

Network IP address: 0.0.0.0 Network mask: 0.0.0.0 Tag: 0 numerical [0-65535]

Adjacent AS number: 0 numerical [0-65535]

BGP Send Policies Already Configured

No.	Type	AM	Network IP address	Network mask	Tag	Adjacent AS no.
1	I	R	0.0.0.0	0.0.0.0	0	0

Buttons: Add, Modify, Delete, Move up, Move down, OK, Save as defaults, Cancel, Help

Figure 87. BGP - Send Policies

Figure 88 on page 144 shows how the BGP originate policies can be defined. The originate policies control which RIP and/or OSPF learned routes are forwarded to BGP neighbors.

ITS0/3746-9x0/BGP - Originate Policies

Configure a BGP Originate Policy

Policy Type: ☐ Inclusive (I) ☒ Exclusive (E)

Address Match (AM): ☐ Exact (E) ☒ Range (R)

Network IP address: 0.0.0.0 Network mask: 0.0.0.0 Tag: 0 num. [0-65535]

BGP Originate Policies Already Configured

No.	Type	AM	Network IP address	Network mask	Tag
1	E	R	0.0.0.0	0.0.0.0	0

Buttons: Add, Modify, Delete, Move up, Move down, OK, Save as defaults, Cancel, Help

Figure 88. BGP - Originate Policies

BGP neighbors can be configured to permit the export of the default route (0.0.0.0) to BGP neighbors. It is not normal to enable this option. It should only be used when the operational consequences are well understood.

An important feature of BGP-4 is route aggregation (see 3.6.2.4, “Route Aggregation” on page 127). Route aggregation minimizes the number of routes advertised and therefore reduces the size of routing tables required in other ASs. Figure 89 shows how BGP aggregated routes can be defined. To prevent BGP from advertising the individual networks that make up the aggregated route, in addition to the aggregated route itself, make sure that (see Figure 87 on page 143) these routes are excluded in the BPG send policy.

ITS0/3746-9x0/BGP - Aggregate Routes

Aggregate Routes

Network IP address: 192.168.0.0 Network mask: 255.255.0.0

No.	Network IP address	Network mask
1	192.168.0.0	255.255.0.0

Buttons: Add, Modify, Delete, Move up, Move down, OK, Cancel, Help

Figure 89. BGP - Aggregate Routes

4.4 Static Routes

Defining static routes on IP routers is in general not recommended because of the administrative overhead, the chances of making errors, and the inflexibility of such a mechanism.

However, static routes are essential when defining a default route (for example, to minimize the number of entries in your routing tables), to define routes to networks accessible via passive routers (that is, not running a routing protocol), and to define routes to networks or ASs where a routing protocol is undesirable for reasons of link cost. The 3746 IP router allows you to define four static routes for each destination. When IP installs a static route into the router table, it chooses the reachable next hop IP address with the lowest cost. IP checks its static routes whenever the reachability of a static route next hop IP address changes:

- If the next hop address currently in the routing table has become unreachable (for example, its network interface has gone down), IP switches the route to the next hop IP address with the lowest cost that is still reachable (among the remaining three).
- If the next hop address currently in the routing table is still reachable but has a higher cost than the next hop IP address that has become reachable, IP switches the route to the reachable next hop IP address with the lowest cost.
- If none of the static route's next hop addresses are reachable, IP deletes the route from the routing table.

So, this means only one next hop IP address at a time is in the routing table: the one with the lowest reachable cost.

Figure 90 on page 146 depicts how to configure static routes. Default, host, network, and/or subnet routes can be defined.

ITS0/3746-9x0/IP Static Routes

Configure a Static Route

☐ Default route

Destination network: 192.168.156.0 Destination mask: 255.255.255.0

Next hop addresses Cost [numerical 1-16]

1	192.168.157.9	2
2		1
3		1
4		1

Static Routes Already Configured

Destination network	Destination mask	Next hop address 1	Cost 1
192.168.156.0	255.255.255.0	192.168.157.9	2

Buttons: Add, Modify, Delete, OK, Cancel, Help

Figure 90. IP Static Routes

A default route is indicated by a network address and a subnet mask of all zeroes (0.0.0.0). A host route is indicated by a non-zero network address and a subnet mask of 255.255.255.255. A network route is indicated by entering the network number appended by zeroes (for example, 9.0.0.0) and a subnet mask reflecting the network class (for example, 255.0.0.0 for class A networks). A subnet route is indicated by entering the subnet number appended by zeroes (for example, 9.132.36.0), with the subnet mask reflecting the subnetting chosen (for example, 255.255.255.0 for class C subnetting).

Note: When RIP is employed as the interior gateway protocol, make sure that a single subnet mask is used within your network.

4.5 Routing Protocols Interoperability

When combinations of routing protocols are used on the IBM 3746-9x0 it is necessary to specify how to pass routes learned by one protocol to another.

Examples of situations where this is required include:

- Advertising static/default route
- Passing route information between RIP and OSPF gateway protocols
- Passing route information between the interior gateway protocol(s) (RIP/OSPF) and BGP

The principle behind route export is simple but there are a number of implementation details that must be understood for the IBM 3746-9x0.

4.5.1 RIP Specifics

The following rules apply when there is both a RIP route and a route from another protocol to the same destination:

- The RIP route will override a BGP route.
- The RIP route will be overridden by an OSPF internal route.
- The RIP route will override a static route if the RIP route's metric is less.
- The RIP route will override a default route if the RIP route's metric is less.
- The RIP route will override an existing OSPF Autonomous System external (ASE) type 1 route if the OSPF external comparison switch is set to type 1, and the RIP metric is lower than the OSPF metric.
- The RIP route will override an existing OSPF Autonomous System external (ASE) type 2 route if the OSPF external comparison switch is set to type 1, or the OSPF external comparison switch is set to type 2, and the RIP metric is lower than the OSPF metric.

ITS0/3746-9x0/RIP - General Parameters

☒ Enable RIP

Originate Default Route

☐ Always originate default route

☐ Originate default route if BGP routes available

From AS number: numerical [1-65535]

To network number:

☐ Originate default route if OSPF routes available

Default route cost: numerical [1-16]

Route Acceptance

Network address:

Add

Modify

Delete

OK Cancel Help

Figure 91. RIP - General Parameters

The 3746 IP router may advertise itself as the default router within the RIP routing domain depending on the values configured within the Originate Default Route field (see Figure 77 on page 132). Therefore you have the option:

- To always originate the default route.
- To originate the default route when a BGP route for a specific destination has been received. The default route will be generated if the route corresponds with a specific AS. AS 0 indicates that the AS number should not be checked.
- To originate the default route when OSPF routes are present in the routing table.

4.5.2 OSPF Specifics

The 3746-9x0 has a configuration option to enable route imports. The type of routes that are imported can be specified. Options exist to import static routes, direct routes, subnet routes, RIP routes, and BGP routes.

All routes are imported as OSPF AS external routes and advertised into the OSPF network using AS external link-state advertisements.

When OSPF imports routes, it is necessary to specify additional parameters to define how the external route will be advertised into the OSPF network. These parameters are the *metric type* (1 or 2) and the *external route tag*. The metric type, if set to 2, defines the metric to the destination to be larger than any internal path. If set to 1 it means that the metric is comparable to internal paths. It is recommended that type 2 be selected to ensure that external routes have a larger metric than internal ones. Type 1 should be selected only if the operational effect of using it is fully understood.

The external route tag provides additional information in the AS external link-state advertisement, although it is currently only defined for use with BGP. The tag, therefore, can normally be set to any value.

ITSO/3746-9x0/OSPF - Imported/Default Routes

☒ Enable routes import

☒ Import static routes ☒ Import direct routes

☒ Import RIP routes ☒ Import subnet routes

☒ Import BGP routes ☒ Enable BGP auto-tag generation

Compare RIP/static routes to OSPF routes: ☐ 1 ☒ 2

Originate Default Route

☐ Always originate default route

☐ Originate default route if BGP routes available

From AS number: 0 num [1-65535]

To network IP address: 0.0.0.0

Originate as type: ☐ 1 ☒ 2

Default route cost: 1 num [1-65535]

Forwarding network IP address: 0.0.0.0

OK Cancel Help

Figure 92. OSPF - Imported/Default Routes

The 3746 IP router may advertise itself as the default router within the OSPF routing domain depending on the values configured within the Originate Default Route field (see Figure 92). Therefore you have the option:

- To always originate the default route.
- To originate the default route when a BGP route for a specific destination has been received. The default route will be generated if the route corresponds with a specific AS. AS 0 indicates that the AS number should not be checked.

The forwarding address indicates to which router default packets are routed. It must be a router interface address that is adjacent (that is, part of the same (sub)net) to the 3746-9x0.

3746 IP routing policies, when multiple routes to the same destination are available, are used in the following order:

- Interior gateway protocols (RIP, OSPF) routes are preferred over BGP routes.
- OSPF routes are always preferred over other types of routes (other than routes to directly attached networks).
 - OSPF intra-area routes are preferred over inter-area routes.
- RIP routes and static routes can override OSPF ASE routes.

It can be configured if external routes will be compared to OSPF ASEs as type 1 or type 2 routes. For details, see “Importing Routes into OSPF” on page 120.

Given the preceding definitions of the configurable OSPF ASE route, static and RIP routes may override OSPF ASE routes in the following situations:

- The OSPF comparison is configured as type 1, and the OSPF ASE route is an OSPF ASE type 2 route.
- The OSPF comparison is configured as type 1, and the OSPF ASE route is an OSPF ASE type 1 route but has a higher cost than the RIP or static route.
- The OSPF comparison is configured as type 2, and the OSPF ASE route is an OSPF ASE type 2 route but has a higher cost than the RIP or static route.

4.5.3 BGP Specifics

The routing information that is sent by BGP to its BGP neighbor depends on the routing information within the IP routing table (that is, locally defined static routes and routes learned from RIP and/or OSPF) and routes received from BGP neighbors. This process is controlled by defining originate and send policies and by defining aggregated routes. For details, see 4.3.2, “BGP Configuration on the IBM 3746-9x0” on page 141.

4.5.4 Static/Default Routes

The default referred to in this subsection is the default route that is generated by BGP if the generate default route option is enabled.

Export default cannot be used to export an intra-AS default route without BGP being configured. If this is a requirement it must be met by the use of the static route (0.0.0.0) exported to the interior gateway protocol within the AS.

The default should be exported to the interior gateway protocol being used within the AS. It will become advertised when and only when BGP communications are established with a neighbor that is configured to generate a default route.

Chapter 5. 3746-9x0 IP Implementation Details

This section gives details of the TCP/IP functions provided by the 3746-9x0 and the way these functions have been implemented. In addition, it gives an overview of the configuration effort required to initialize the 3746 for IP transport. This section should be read in conjunction with Chapter 4, "IBM 3746-9x0 and IP Routing Protocols" on page 131.

5.1 Multiple TCP/IP Stacks

To understand how the TCP/IP functions have been implemented the reader should know of which components the 3746-9x0 is comprised (see Figure 93).




Figure 93. 3746-9x0 IP Base

For the 3746 IP functions, the following components should be identified:

- The 3746-9x0 controller base
- The network node processor
- The service LAN
- The service processor

The 3746-9x0 controller consists of multiple independent adapters and a connectivity switch (CS) that enables data transfer between the adapters. Each adapter contains a single processor and one or more couplers.

Adapters are identified by the function performed: token-ring adapter (TRA), communication line adapter (CLA), ESCON adapter (ESCA), and controller bus and service adapter (CBSA). The number and type of adapters installed is customer-dependent. However, each 3746-9x0 must contain a CBSA.

Each 3746-9x0 processor is a general purpose processor operating independently, having its own storage. Each processor is identified by the function it performs: token-ring processor (TRP), communication line processor (CLP), ESCON processor (ESCP), and controller bus and service processor (CBSP). The number and type of processors installed is customer-dependent. However, each 3746-9x0 must contain a CBSP. Tasks executed on the 3746-9x0 processors are the data link control (DLC) functions and all higher layer functions needed to perform the APPN, frame relay, and IP functions.

The couplers provide the electrical and mechanical interface functions to enable communication equipment to connect to the 3746-9x0. The couplers are identified by their function: token-ring interface coupler (TIC), line interface coupler (LIC), and ESCON coupler (ESCC). The number and type of couplers installed is customer-dependent. However, each 3746-9x0 must contain a TIC connected to the CBSP.

The mandatory token-ring coupler controlled by the CBSP (port 2080) attaches the 3746-9x0 base to the service LAN. This 16-Mbps token-ring LAN enables communication between the 3746-9x0 base, the network node processor and the service processor.

Figure 94 depicts how TCP/IP functions have been implemented on the 3746-9x0 components identified.




Figure 94. 3746-9x0 IP Base

Each 3746-9x0 processor, including the CBSP, provides IP and data link control (DLC) functions. The DLC functions enable data transport to adjacent nodes. The

DLC functions are shared between protocol stacks (IP, APPN/HPR, NCP/Subarea, and frame switching). The IP functions allow IP transport between each of the 3746-9x0 adapters and externally attached IP equipment. The code that performs the IP routing functions is present on all processors. However, it will only be activated on processors on which an IP interface has been configured. To perform the IP routing, each adapter maintains a routing table cache. The routing cache is built during the transport of IP data, using routing information learned from the CBSP.

In addition to the IP forwarding and DLC functions, more functions are required. Address resolution protocols (ARP and InARP) run on each processor, while the dynamic IP routing protocols (RIP, OSPF, and BGP) run on the CBSP. The CBSP is responsible for processing routing protocol packets (RIP/OSPF/BGP) received from, and generating routing protocol packets destined for, external equipment. The adapters are instrumental in forwarding the packets between the external equipment and the CBSP.

The CBSP builds and maintains a routing table using static routes and IP interfaces defined on the 3746-9x0 and routing information learned from the dynamic routing protocols. To minimize the routing information maintained on each of the adapters, each processor maintains a routing table cache that contains a subset of the information stored in the CBSP routing table.

Attached to the service LAN is the network processor (NNP). The main functions of the NNP are the 3746 IP operator and the SNMP functions. The NNP is not involved in the IP packet forwarding process in any way. To access the 3746 IP operator functions one can TELNET into the NNP. In addition, the NNP stores the 3746 IP configuration files on its hard disk. When the 3746 IP code is restarted, each adapter retrieves a copy of the 3746 IP configuration file and configures itself. For details about the operator functions, see 5.10, "3746-9x0 IP Management" on page 214, and for details about the SNMP functions, see 5.11, "3746-9x0 SNMP Functions" on page 225.

Also attached to the service LAN is the service processor (SP). The SP provides a repository of the microcode running on the 3746-9x0 adapters. Each 3746-9x0 adapter loads its microcode from the SP during IML. The 3746-9x0 controller configuration and management (CCM) tool can be accessed from the SP.

5.2 Route Caching

During IP forwarding, if a 3746-9x0 processor is not capable of making a routing decision based on the information maintained in its routing table cache, the appropriate routing information is retrieved by querying the CBSP. Host routes learned from the CBSP are stored within the routing table cache and are used to route successive IP datagrams. The cache contains one entry per destination host and not per destination network, this is to prevent less specific routes from masking more specific routes.

Cached entries will remain active until a change occurs in the CBSP routing table. After a change in the CBSP routing table, the CBSP will initiate a purge of the routing caches on all other processors. Changes in the CBSP routing table result from new routes learned via the dynamic routing protocols and will not occur during a steady-state. Cached entries are subject to an aging mechanism and will be removed periodically.

Each route a 3746-9x0 processor learns from the CBSP can be any of the following:

- Host route** A route for a specific host IP address
- Direct route** A route for IP networks locally attached to the 3746-9x0
- (Sub)network route** A route for a (non-local) (sub)network
- Default route** A defined route to be used when no explicit route exists

As routers make routing decisions based on the *destination* IP address in the datagram to be forwarded, the route the CBSP returns will be the route most closely *matching* the destination IP address. A host route is considered the most and the default route the least-specific route.

5.2.1 Distributed Routing and ARP Caching

The following section gives a brief explanation of how the routing decision is distributed between the CBSP and processors in the 3746. It also explains how ARP caching is done. Only the caches and information needed for this explanation are shown.




Figure 95. Distributed Routing #1

In Figure 95 the routing and ARP caches are empty on each processor. The CBSP has various routes which may have been obtained from dynamic routing protocols or static definitions.




Figure 96. Distributed Routing #2

In Figure 96 an IP datagram for destination IP address 9.10.1.1 arrives at port 2304 (1). There is no route for this destination in the processor cache, therefore the processor queries the CBSP with a GET_ROUTE request (2) for that IP destination.




Figure 97. Distributed Routing #3

In Figure 97 a GET_ROUTE response (3) is sent by the CBSP indicating a next hop IP address of 10.1.1.2, which can be reached via physical port 2176.

The processor forwards the IP datagram to the processor for port 2176 (4).




Figure 98. Distributed Routing #4

In Figure 98 processor 2176 receives the datagram for destination IP address 9.10.1.1 with a next hop of 10.1.1.2, but has no MAC address for that IP address in its ARP cache.

The processor issues an ARP request for 10.1.1.2 (5), and the router 10.1.1.2 replies (6) with its MAC address 400000000002.

The processor can now forward the IP datagram onto the token-ring with the destination MAC address (7), which is also added to the ARP cache.

5.2.2 Default Route Caching

All types of routing entries learned from the CBSP are cached on the 3746-9x0 processors except the default route. Figure 99 on page 158 and the following text explains why the default route is never cached.




Figure 99. Default Route Caching

The following list explains Figure 99:

1. An IP datagram arrives at port 2304, with a destination IP address of 9.132.99.1.
2. The processor has no route to the IP destination in its cache, so it queries the CBSP for a route. In the CBSP routing tables, there is no matching route so the default route will be returned. The default route points to a next hop of 9.132.8.4, which is via port 2208.
3. An IP datagram arrives at port 2304, with a destination IP address of 9.132.57.8. If the default route were cached, then it would be the best match in the processor cache, and this datagram would be routed to port 2208. To avoid this, the default route (received in 2) is cached as a host route to 9.132.99.1. This means that the processor has no route to 9.132.57.8, and therefore queries the CBSP for a route.
4. The CBSP has a route with next hop 9.132.8.57 via port 2176, this route is returned to the processor, used to route the packet, and cached for further use.

A drawback of not being able to cache the default route is that the processor queries the CBSP for every datagram for which the default route applies. To minimize this processor-to-CBSP overhead, 3746-9x0 processors add a host route for each destination that has been routed using the default route. As can be seen in the previous example, routing of 9.132.99.1 has resulted in a cached destination host entry on the processor.

Route caching reduces the processor-to-CBSP query overhead, but may result in large routing tables. Make sure that for your environment the size of the routing table and number of destination addresses that can be cached are appropriately set. For default and maximum values, see Figure 100 on page 159. Note that in many cases the use of dynamic routing protocols or static routes avoids the need of default routing and, therefore, caching destination addresses.

Figure 100. IP General Parameters

5.2.3 Displaying the CBSP Routing Tables

During problem determination it is often useful to display the 3746 IP routing information. Be aware of the distinction between the routing caches maintained on the 3746-9x0 processors and the full routing table kept on the CBSP. IP datagrams are routed using a processor routing cache, unless no route is available and the CBSP is queried. Routing caches are subsets of the CBSP routing table with destination host routes added for datagrams routed via the default route (see the discussion in the previous section). 6.3.3, “Dump CBSP Routing Table Display” on page 268 and 6.3.4, “3746-9x0 Processor IP Routing Cache Display” on page 268 discuss how to display the CBSP and other 3746-9x0 processor routing cache information using the NNP TELNET interface.

Type	Dest net	Mask	Cost	Age	Next hop[s]
Sbnt	9.0.0.0	FF000000	1	429471580	None
SPF*	9.24.104.0	FFFFFF00	1	429471598	TKR/1
Dir*	192.168.150.0	FFFFFF00	1	429471598	TKR/1
Dir*	192.168.152.0	FFFFFF00	1	429471580	TKR/0
SPF	192.168.152.3	FFFFFFF	0	429471598	SINK/0

Routing table size: 768 nets (46080 bytes), 5 nets known

Figure 101. CCM Routing Tables Display

The top line of the display shows the device address as 2080, which is the CBSP.

The rightmost column shows the physical interface identifiers, which are the next hop within the 3746, corresponding to the IP net number (second column from left). The 3746 hardware addresses (port numbers) corresponding to these identifiers can be displayed by using the telnet command shown in Figure 227 on page 270, or from CCM by selecting **Management -> IP Specifics -> Configuration**.

5.3 3746 IP Interfaces Overview

The 3746 IP router supports attachment of IP equipment using ESCON, token-ring, Ethernet, or serial lines. For serial line-attached equipment, frame relay, X.25, and point-to-point protocol (PPP) connections can be used. Currently only leased line PPP connections are supported.

Each of these interfaces are discussed in more detail in the following sections.

5.4 ESCON

ESCON attachments can be used to provide native IP transport, using the channel data link control (CDLC), between the 3746 IP and host system(s) running TCP/IP for MVS ($\geq$ Version 3.1), or OS/390 $\geq$ Version 2.4. The host systems can be directly attached to the 3746-9x0 or connected via an ESCON director (ESCD).

The ESCON implementation on the 3746-9x0 is based on the concept of host links and link stations. A host link is a logical connection between a host system and the 3746-9x0. A link station represents a logical point-to-point connection between the 3746-9x0 and host programs such as TCP/IP for MVS, VTAM, and/or TPF. Host link and link station definitions need to be entered before IP (or, for that matter, APPN or SNA subarea) communication is possible.




Figure 102. ESCON Port, Host Link, and Link Station

In each 3746-9x0 you can install up to 16 ESCON adapters (15 on a 3746-900 attached to a dual CCU Model 3745). Each ESCON adapter controls a single ESCON coupler. An ESCON coupler provides a single physical connection to a single host system, unless you are using an ESCON director (ESCD). An ESCD provides a physical connection to multiple host systems.

On each ESCON adapter you can define up to 16 host links. To each host system you can define only one host link, unless the host system is using the ESCON Multiple Image Facility (EMIF), in which case you can define a separate host link (up to 16) to each of the logical partitions (LPARS). On each host link you can define one or multiple link stations. However, the maximum number of link stations that can be defined per ESCON adapter is 16.




Figure 103. Multiple 3746 IP ESCON Links

Note: Using the above figures, a single 3746-9x0 can provide an ESCON attachment to up to 256 (16*16) host systems and communicate with up to 256 host programs.

The ESCON configuration process consists of four phases:

1. Definition of the physical ESCON fiber and its connectivity (see 5.4.3, “ESCON - Fiber and Host Link Configuration” on page 165).
2. Definition of one or more host links over ESCON fibers (see 5.4.3, “ESCON - Fiber and Host Link Configuration” on page 165).
3. Definition of up to 16 link stations on each host link
4. Assigning the link stations to either the APPN, IP, or NCP protocol stacks

Figure 103 depicts an example in which a number of link stations have been assigned to the 3746 IP function. Each IP connection is a separate point-to-point connection to which an IP address and subnet mask has to be assigned at both ends.

5.4.1 ESCON Port Sharing

As discussed in the previous section up to 16 link stations can be defined per ESCON adapter. Each link station provides a point-to-point connection and is the logical equivalent of a port.




Figure 104. ESCON Port Sharing

ESCON ports can be defined and simultaneously activated from all protocol stacks available (IP, APPN, and (3746-900 only) NCPs running on CCU-A and CCU-B on the attached 3745). On a 3746-900, each link station must be assigned to either 3746 IP, NCP/CCU-A, NCP/CCU-B, or the 3746-900 NN function; on a 3746-950, a link station must be assigned to either the 3746-950 IP or the 3746-950 APPN/DLUR function.

Link stations map one-to-one to I/O addresses defined on attached host systems. The I/O addresses are used by host programs to communicate with the corresponding protocol stack. The 3746-9x0 IP and NCP IP protocol stacks can only communicate with TCP/IP for MVS ($\geq$ V3.1) and OS/390 ($\geq$ V2.4), while the 3746-9x0 APPN and/or NCP SNA protocol stacks can communicate with VTAM and/or TPF.

Link stations are independently assigned; therefore, an ESCON port can be shared between up to four (two on the 3746-950) 3746-9x0 protocol stacks.

5.4.2 CDLC Configuration Rules

There are some basic rules that must be observed when specifying IP addresses for 3746 to TCP/IP for MVS CDLC connections. Each IP CDLC connection is a point-to-point subnetwork, therefore in some cases (which are explained later), the same IP addresses may be used multiple times. On the other hand there are conditions where the same addresses may not be used more than once, or where different subnetwork addresses must be used.

The rules are as follows:

- General Rules:
 - One MVS host system can run multiple IP address spaces. Each IP address space is a full TCP/IP for MVS protocol stack.
 - Each IP address space is independent of the others on the same host.
 - Each IP address space can run multiple CDLC instances.
 - Each CDLC instance of a given IP address space must run on a separate ESCON fiber.
- TCP/IP for MVS Rules:
 - Each CDLC instance of a single IP address space must have a different next-hop address.
 - CDLC instances of different IP address spaces may have the same next-hop address.
 - To enable routing between MVS systems via the 3746, either a different subnet must be used for each MVS, or static routes would need to be defined in each MVS.
- 3746 Rules:
 - For hosts with only a single IP address space and a single CDLC instance in that address space users can define a single subnet between the hosts and 3746 stations.

 The 3746 may have the same next-hop address on all CDLC connections as each CDLC instance is separate from the others. But each host CDLC instance must have a different address so that the 3746 can determine over which interface to route packets destined for that address.
 - For hosts with multiple IP address spaces and a single CDLC instance in each address space: users can define a single subnet between the hosts and 3746 stations.

 As before, the two MVS instances are totally independent of each other; therefore they may route to the same next-hop address in the 3746.




Figure 105. CDLC IP Addressing Example

- For hosts with multiple IP address spaces and multiple CDLC instances in each address space users must define multiple subnets between the hosts and 3746 stations. Two CDLC instances of the same IP address space may not be in the same subnet.




Figure 106. CDLC IP Addressing Example

5.4.3 ESCON - Fiber and Host Link Configuration

ESCON support for host links and links stations is dependent upon various factors in the MVS host, and the ESCON equipment used. Each host link supports a maximum of 16 link stations.

A host link goes to a single host in *basic mode*, or to a single logical partition (LPAR) when in *partitioned mode*. A host link never goes to multiple hosts or LPARs.

Multiple host links may be supported over a single ESCON fiber. A host running in partitioned mode without EMIF requires an ESCON fiber to each LPAR. With EMIF installed, a single ESCON fiber can support host links to multiple LPARS.

Figure 107 shows two configurations. When the host is in basic mode, then a single host link is defined to the host. This host link supports multiple link stations which can be use by various protocol stacks (VTAM, TPF, and IP for example) running under the MVS on the host.

If the host is running in partitioned mode and ESCON Multiple Image Facility (EMIF) is not supported, then a single host link across a single ESCON fiber can only be assigned to one LPAR. Protocol stacks in this LPAR can use the host link.




Figure 107. ESCON to MVS Connectivity without EMIF

Figure 108 on page 166 shows two configurations. When the host is in partitioned mode without EMIF, then two ESCON fibers are needed between the 3746 and host to provide connectivity to both LPARs. With EMIF installed, then a single ESCON fiber can support host links going to multiple LPARs, but each host link still only attaches to a single LPAR.




Figure 108. ESCON to MVS Connectivity

Figure 109 shows configurations with a host in basic mode, and one in partitioned mode with EMIF. The restrictions previously mentioned on the MVS side apply, but as an ESCON director is between the 3746 and host systems, a single ESCON fiber from the 3746 to ESCON director can support all the host links needed.




Figure 109. ESCON to MVS Connectivity

5.4.4 ESCON - Link Stations and IP Configuration

The ESCON link station definition and IP configuration for the link station, which follows the CCM definition of an ESCON port and host link, is a very straight forward process consisting of:

- Creating a link station
- Assigning this station to a protocol stack (in this case IP)
- Assigning an IP address and subnet mask to the link station
- Enabling a dynamic routing protocol on this interface (or define static routes)

Assigning an IP address and subnet mask (see Figure 110) is mandatory. If necessary, the default channel adapter slowdown time (CASDL), the attention timer (TIMEOUT), and the delay time (DELAY) timer can be changed by clicking on the DLC parameters.

ESCON Station Configuration - Port 2240

APPN host Link: IP host Link: P2240HL Number of host Links: 1

Example: Stations

Host or Partitions	VTAM/TPF	VTAM/TPF
Unit Add:	Unit Add:	Unit Add:
01 02	03	

3746-900

3745

CCU-B

Add:01

Add:03

CCU-A

Add:02

Configure an ESCON Station

Network? ☐ APPN (A) ☒ IP (I) ☐ SNA/Subarea (S)

UTAM ☐ TPF

Name: S2240IP8

PU type: ☒ 1 ☐ 2.1 ☐ 5

Unit Address (UA): 8 hex

IPL through that station? ☐ Yes ☒ No

On which CCU? ☐ CCU-A ☒ CCU-B

IP address: 192.168.202.8

IP subnet mask: 255.255.255.0

Comments (optional):

ESCON Stations Already Configured

Name	Network	PU	UA	CCU
S2240IP8	IP	1	8	No
S2240IP9	IP	1	9	No

Buttons: Add, Modify, Delete, DLC parameters..., APPN parameters..., OK, Cancel, Help

Figure 110. ESCON Station - APPN/IP DLC Parameters

The unit address of the link station must be matched to the device number in the host system IOCP definitions. It is important to remember that for the 3746 the 16 link stations are numbered 0x01 to 0x10, whereas in IOCP devices are numbered 0x00 to 0x0F.

Enabling dynamic routing protocols is optional. It is only required when the TCP/IP for MVS system has multiple IP connections, and static routing does not suffice. The only dynamic routing protocol currently supported by TCP/IP for MVS Version 3.1 is RIP. Configuring RIP is detailed in 4.1.1, "RIP Configuration on the IBM 3746-9x0" on page 132. For details on configuring static routes, see 4.4, "Static Routes" on page 145. See 7.6.1, "MVS Definitions" on page 306 for an example of the required host definitions.

5.4.5 ESCON - IP Configuration Example

The following diagrams (Figure 113 on page 169, Figure 114 on page 170, and Figure 115 on page 171) show an example ESCON configuration. In this example, a single host link is configured from a 3746 to a host running in basic mode. The APPN and IP protocol stacks share this link, with a single link each over the host link.

Figure 111 shows the IOCP definitions needed for the 3746.

```
*IOCP*****
*IOCP    3745/3746 THROUGH SWITCH E1          *
*IOCP*****
          CNTLUNIT CUNUMBR=920,UNIT=3745,PATH=(18),      +
          CUADD=0, LINK=(E0),UNITADD=(01,32))
DEV920    IODEVICE UNIT=3745,ADDRESS=(920,32),CUNUMBR=(920),UNITADD=01
```

Figure 111. IOCP Definitions for 3746

IOCP definitions for 32 devices starting at address 920 were created for this 3746. The IP station is station number 16 which is device number 92F. Figure 112 shows the TCP/IP for MVS definitions needed for the 3746.

```
*****
*          3746 Station 16 on 920          *
*****

DEVICE DEVXA8C  CDLC 92F 200 200 2060 2060
LINK  LINKXA8C CDLC 0    DEVXA8C

HOME  192.168.202.18    LINKXA8C

START  DEVXA8C
```

Figure 112. TCP/IP for MVS Definitions for 3746 IP

On the ESCON port configuration screen, the port is created and we specify that APPN and IP will use it. The ESCON director definitions are also made for the ESCON link.




Figure 113. ESCON Port Configuration Example

On the ESCON host links configuration screen, the host link is defined. As we are connecting to a single host in basic mode, a single host link is needed in this case. CHPID and host link addresses are specified for this host link.




Figure 114. ESCON Host Link Configuration Example

The ESCON station configuration screen allows us to define our APPN station (Unit Address=0x09) and IP station (Unit Address=0x10). Here we also define the local IP address for our IP station.




Figure 115. ESCON Station Configuration Example

Note: The IOCP definitions shown are one possible way of defining the 3746. Another possibility would be to use a single *device* statement for each address, in which case UNITADD=10 would be used for this station.

5.5 Token-Ring

Token-ring attachments can be used to provide IP transport between the 3746 IP and IP stations or routers running TCP/IP.

For the IP communication over token-ring, the 3746-9x0 uses (connectionless) IEEE 802.2 logical link control (LLC). IEEE 802.2 LLC requires the use of a local and a remote service access point (SAP). For IP communication both local and remote SAP should be equal to X'AA', indicating SNAP encapsulation (see Figure 116). The (optional) routing information field (RIF) in the MAC header indicates the support for source-route bridging, which is available by default.




Figure 116. Token-Ring IP Encapsulation

Per 3746-9x0 you can install up to 16 token-ring adapters (15 on a 3746-900 attached to a dual CCU Model 3745). Each token-ring adapter can control up to two token-ring interface couplers (TIC3). Each token-ring coupler provides a single physical connection to a token-ring LAN.

Note: Using the above figures and including the TIC3 controlled by the CBSP, the 3746-9x0 supports the attachment of up to 33 (16*2 + 1) token-ring LANs (31 on a 3746-900 attached to a dual CCU Model 3745).

The token-ring configuration process consists of two phases: a general definition part in which token-ring port specifics are defined and the actual IP port configuration.

5.5.1 Token-Ring Port Sharing

Token-ring ports can be defined on all protocol stacks. However, simultaneous activation of the port is limited to the IP protocol stack, APPN protocol, and (for 3746-900 only) either the NCP running on CCU-A or the NCP running on CCU-B of the attached 3745. Therefore, on a 3746-900, a token-ring port can be shared between the 3746 IP, a single NCP and the 3746-900 NN function; on a 3746-950, the token-ring port can be shared between the 3746-950 IP and the 3746-950 NN function. The 3746-9x0 data link control layer differentiates between traffic to any of these protocol stacks using the service access points (SAPs) within the token-ring MAC header.

ACF/NCP V6R2 introduced NCP support for TIC3 sharing with 3746 NNP or 3746 IP, while 3746 EC Level D46100A (ECA155) introduced 3746 support for TIC3 sharing.




Figure 117. Token-Ring Port Sharing


When receiving data the 3746-9x0 will accept the following destination SAPs:

- X'AA' - Data to/from 3746 IP stack
- X'04' - SNA data to/from NCP protocol stack (3746-900 only)
- X'C8' - HPR non-ERP data to/from NCP protocol stack (3746-900 only)

The adjacent node can negotiate a different HPR SAP if so desired. The 3746 will always default to X'C8' if no other SAP is negotiated.

5.5.2 Token-Ring - General Configuration

To define a token-ring port you have to specify (see Figure 118) its IP and APPN (if relevant) names, the speed of the token-ring LAN, and the local MAC address. The APPN and HPR SAPs are only relevant if the port is also used for APPN traffic (see 5.5.1, “Token-Ring Port Sharing” on page 173).



The image shows a 'Token-Ring Port Configuration' dialog box. It has a title bar with a minus sign and the text 'Token-Ring Port Configuration'. Inside, there's a section 'Configure Token-Ring Port 2144'. Below this, there are two radio buttons: 'APPN' (unchecked) and 'IP' (checked). To the right of the 'IP' button is the text 'Number of APPN stations configured: 0'. Below the radio buttons are two text input fields: 'APPN name' (empty) and 'IP name' (containing 'P2144IP'). To the right of the 'IP name' field is a 'Speed' section with two radio buttons: '4' (selected) and '16', followed by 'Mbps'. Below these are three more text input fields: 'Local MAC address (LAA or UAA):' (containing '40002144061A'), 'APPN local SAP (LSAP):' (empty), and 'IP maximum transmission unit:' (containing '4060'). To the right of the MAC address field is the word 'hexadecimal'. To the right of the LSAP field is 'hexadecimal [04-9C]'. To the right of the MTU field is 'bytes [516 - 4060]'. Below these fields is a 'Comments (optional):' text area containing 'Local Office LAN'. At the bottom of the dialog, there are five buttons: 'APPN parameters...', 'Connection network...', 'DLC defaults...', 'IP parameters...', and 'OK'. Below these are four more buttons: 'Delete', 'Stations...', 'Cancel', and 'Help'.

Figure 118. Token-Ring Port Configuration

Make sure the IP maximum transmission unit (MTU) is high enough to prevent fragmentation of IP datagrams on the 3746 IP router (see 2.8, “Fragmentation and Reassembly” on page 80). Select IP parameters to enter the IP over token-ring configuration. For details, see 5.5.3, “Token-Ring - IP Configuration.” For IP transport no station needs to be defined.

5.5.3 Token-Ring - IP Configuration

After a token-ring port has been added and/or enabled for IP, one has to specify the IP configuration. The IP configuration for token-ring ports is a very straight forward process consisting of:

- Assigning an IP address and subnet mask to the token-ring port

The 3746 IP router allows you to specify multiple IP addresses and subnet masks to the same physical interface. When defining multiple IP addresses, the 3746 IP router connects to multiple subnets. Note that although the 3746-9x0 is capable of routing between subnets on the same physical interface, in general this is not recommended. See also the discussion in 2.7.1.2, “Multiple Logical Subnets Per Physical Network” on page 76.

- Enabling dynamic routing protocol(s) for each IP address, or possibly defining static routes

The 3746 IP dynamic routing protocols operate independently for each IP address defined on a token-ring interface; therefore configuration effort is required for each IP address defined.

Assigning an IP address and subnet mask (see Figure 119 on page 175) is mandatory. At least one, but optionally multiple (up to 16) IP addresses, can be defined. If multiple IP addresses are defined, make sure that the addresses are part of different subnets. If multiple IP addresses in the same subnet are assigned, only the first one will be used.

IP Over Token-Ring Parameters							
Port: 2144	Name: P2144IP						
General Automatic reactivation? ☒ Yes ☐ No ☒ Enable source routing RIF timer: 120 seconds [0-4096]							
Addresses <table border="1"> <thead> <tr> <th>IP address:</th> <th>Subnet mask:</th> </tr> </thead> <tbody> <tr> <td>9.24.104.168</td> <td>255.255.255.0</td> </tr> <tr> <td>192.168.150.3</td> <td>255.255.255.0</td> </tr> </tbody> </table>		IP address:	Subnet mask:	9.24.104.168	255.255.255.0	192.168.150.3	255.255.255.0
IP address:	Subnet mask:						
9.24.104.168	255.255.255.0						
192.168.150.3	255.255.255.0						
Add Modify Delete							
OK Save as defaults Cancel Help							

Figure 119. Token-Ring Port - IP over Token-Ring Parameters

The RIF timer indicates how long (in seconds) the router holds information stored in the routing information field (RIF), before it is aged out.

Enabling dynamic routing protocols is optional. It is only required when any of the token-ring attached devices have multiple IP interfaces, and static routing does not suffice. The dynamic routing protocols supported by the 3746-9x0 are RIP, OSPF, and BGP. Configuring RIP is detailed in 4.1.1, “RIP Configuration on the IBM 3746-9x0” on page 132, configuring OSPF is detailed in 4.2.5, “OSPF Configuration on the IBM 3746-9x0” on page 136, and configuring BGP is detailed in 4.3.2, “BGP Configuration on the IBM 3746-9x0” on page 141. For details on configuring static routes, see 4.4, “Static Routes” on page 145.

5.6 Ethernet

As discussed in Chapter 1, “3746 IP Router - Overview” on page 1, 3746 IP support for Ethernet-attached equipment can be realized by ordering one of the 3746-9x0 Ethernet features #5631 or #5632.

Both features (see Figure 15 on page 17) provide Ethernet support by physically connecting the 3746-9x0 to a token-ring LAN and, using an IBM 8229 translational bridge, bridging the token-ring LAN to an Ethernet LAN.

5.6.1 Ethernet Configuration

Because of the translational bridging, token-ring rather than Ethernet configuration is required to enable Ethernet access to the 3746-9x0. For details on how to configure the 3746-9x0 for token-ring, see 5.5.3, “Token-Ring - IP Configuration” on page 174.

5.7 Frame Relay

To support IP transfer over serial connections, the 3746 IP router offers three solutions: frame relay, X.25, and point-to-point (PPP) links.

Using frame relay to transport IP has the following advantages:

- High throughput and low delay

Utilizing the core aspects (error detection, addressing, and synchronization) of the Link Access Protocol D-Channel (LAPD) data link protocol, frame relay eliminates all network layer (Layer 3) processing. By using only the core aspects, frame relay reduces the delay of processing each frame.

- Traffic scheduling

For each PVC a COMRATE or CIR value will be assigned to 3746 IP. The COMRATE parameter operates in conjunction with the DATABLK parameter set for that frame relay line. When output queuing takes place, the 3746 frame relay transmit scheduling mechanism will assign each PVC a fair share of the bandwidth by transmitting a maximum of COMRATE times DATABLK bytes before servicing the next PVC. See also the discussion in 5.7.4, "Frame Relay Scheduling" on page 178. In addition, the Bandwidth Reservation System (BRS) can be used to divide the PVC bandwidth up between multiple protocols that are sharing a PVC.

- Circuit access and control

As the 3746 IP router can dynamically learn about the availability of non-configured circuits, you can control access to those new circuits.

- Network management option

As your network requires, the frame relay protocol can operate with or without a local network management interface.

- Multiplexing protocols

They use one PVC to pass multiple protocols.

Frame relay provides no error correction or retransmission functionality. To provide error-free end-to-end transmission of data, frame relay relies on the intelligence of the host devices.

The 3746-9x0 IP frame relay support uses RFC 1490 encapsulation to transfer IP data over frame relay. IP over frame relay connections are supported to any router complying with RFC 1490. (Examples are NCP V7R3 and above, IBM 2210, IBM 2216, and IBM 6611.)

On each 3746 IP frame relay interface one or multiple permanent virtual connections (PVCs) can be established. Each PVC represents a point-to-point connection. PVCs used by the 3746-9x0 IP router can be defined statically or can be learned dynamically. Dynamic (not predefined) PVCs can only be used on frame relay connections on which the frame relay local management interface (LMI) procedures have been activated.

On a single frame relay line, multiple (local) IP addresses and subnet masks can be defined. Note that IP addresses are defined per frame relay line rather than per PVC. When a subnet is defined over several PVCs on the same port, the subnet must be fully meshed; that is, all the routers of this subnet must be interconnected.

The IP addresses of remote IP routers can be predefined or learned dynamically. To learn the remote IP address, the 3746 IP uses the Inverse ARP (InARP) protocol.

5.7.1 Orphan Circuits

An orphan circuit is any PVC that is not configured but is learned indirectly through the actions of the LMI. The Inverse ARP protocol for frame relay maps the orphan circuit DLCI (a hardware address) to an IP address. To do this the 3746 IP router sends an Inverse ARP request over the orphan circuit to the remote router, which returns its IP address.

Orphan circuits are treated the same as configured circuits, except that you may enable or disable their use with the enable and disable commands.

By disabling orphan circuits, you add a measure of security to your network by preventing any unauthorized entry into your network from a non-configured circuit. By enabling orphans, you allow the router to forward packets over circuits you did not configure. When using orphan circuits, bandwidth management or BRS parameters for example cannot be predefined.

5.7.2 Multicast Emulation

Multicast emulation is an optional feature that allows protocols requiring multicast to function properly over the frame relay interface. With multicast emulation, a frame is transmitted on each active PVC. By using the enable and disable multicast commands, you can turn this feature on or off.

5.7.3 Frame Relay Network Management

The supplier of the frame relay network backbone provides frame relay network management. This provides frame relay endstations (routers) with status and configuration information concerning PVCs available at the physical interface.

The 3746-9x0 supports the ANSI Annex D management and ITU-T Q.933 Annex A management entities. You can turn these entities on or off. Specifically, frame relay network management provides the following information:

- Notification of additional PVCs (orphans) and whether they are active or inactive or of any PVC deletions.
- Notification of PVC status separate from a router's polled status request.
- Notification of flow control through the FECN and BECN bit settings.
- Notification of the availability of a configured PVC. The availability of a PVC is indirectly related to the successful participation of the PVC endpoint in the heartbeat polling process.
- Verification of the integrity of the physical link between the endstation and network by using a keep-alive sequence number interchange.
- Inclusion of CIR as part of the PVC status information.

Although the frame relay interface supports both types of network management, it is not mandatory to activate the LMI procedures. For example, you may want to disable management for back-to-back testing.

5.7.4 Frame Relay Scheduling

For frame relay there are two different levels where scheduling occurs. The first is at the DLCI level; the second is at the protocols level. For the DLCI level, we can use COMRATE or CIR to control the transmission of data into the network. At the protocol level, we can use BRS to assign portions of a PVC's bandwidth to different protocols.

5.7.4.1 Committed Information Rate (CIR) and Burst Sizes

The 3746-9X0 provides CIR support which is *not* available under NCP. The maximum number of bits per seconds that an endstation can transmit into the network is bounded by the *access rate* of the user network interface. The access rate is limited by the line speed of the user network connection and established by subscription.

The maximum committed amount of data that a user may offer to the network is defined as *committed burst size* (B_c). B_c is a measure for the volume of data for which the network will guarantee message delivery under normal conditions. It is measured during the *committed rate measurement interval* (T_c).

Endstations are allowed to transmit data in excess of the committed burst rate. The *excess burst size* (B_e) has been defined as the allowed amount of data by which a user can exceed B_c during the committed measurement rate interval T_c . If spare capacity exists, the network will forward the data to its destination. However, the network is free to mark the data as discard eligible (DE).

The *committed information rate* (CIR) has been defined as the allowed amount of data that the network is committed to transfer under normal conditions. The rate is averaged over an increment of time T_c . The CIR is also referred to as *minimum acceptable throughput*. Remember that CIR is implemented only for 3746 controlled APPN and IP lines.

B_c and B_e are expressed in bits, T_c in tenths of seconds, while the access rate and CIR are expressed in bits per second. CCM allows you to define the outgoing B_c , B_e , and T_c values for each DLCI, from which the CIR is then computed. The access rate is valid for each network interface.

For B_c , B_e and CIR incoming and outgoing values can be distinguished. If the connection is symmetrical the values in both directions are the same. For permanent virtual circuits B_c (incoming and outgoing), B_e (incoming and outgoing) and CIR (incoming and outgoing) are defined at subscription time. They are negotiated for SVCs at call establishment time. T_c is calculated as depicted in Table 5 on page 179.

Table 5. Measurement Interval Calculation			
CIR	B _c	B _e	Measurement Interval (T _c)
> 0	> 0	> 0	T _c = B _c /CIR
> 0	> 0	0	T _c = B _c /CIR
0	0	> 0	T _c = (B _e /Access Rate) ²
Notes: 1. This table depicts the valid parameter configurations. Other configurations are under further study. 2. When the two communicating endstations have different access rates, the network may define a smaller T_c value.			

Individual CIRs on a physical connection are always lower than the access rate; however, the sum of CIRs defined can be larger than the access rate. An example could be a network connection with an access rate of 256 kbps on which three virtual circuit have been defined: two having a CIR of 128 kbps each, and one having a CIR of 64 kbps.

Optimal values for the above parameters depend on network implementation, availability of spare network capacity, charging methods, type of user device and performance required. These are only a few considerations, and careful study of the total network is required (as well as the more immediate future changes that are anticipated).

Networks may mark frames above B_c with discard eligible (DE) but have plenty of spare capacity to transport the frame, or may instead have limited capacity and discard excessive frames immediately. Networks may mark frames above B_c+B_e with discard eligible (DE), and possibly transport it, or just drop the frames as suggested by ITU-T I.370.

The network manager always tries to balance costs and performance, and examines the frame relay service provider charging schemes. Networks may implement fixed charging depending on access rate, a scheme dependent on CIR, B_c and B_e, or more sophisticated schemes, for example charging on number of bits transported and charging progressively for data above B_c or B_c+B_e. Depending on the charging scheme employed, subscribing to high values of CIR, B_c and B_e, may lead to high networking costs. It should be examined if the performance gain, if any, counterbalances the additional networking expenses.

Many devices have limited control over the volumes of data they send into the network. Assuming that flow control mechanisms implemented on top of the layer 2 core function are not inhibiting data transfer, data will be transmitted with a speed up to the network access rate. If the device has only one DLCI active, or has (temporarily) data to send for one DLCI only, the data rate on a single DLCI may be equal to the network access rate. If the sum of committed and excess burst size (B_c+B_e) is lower than the access rate times T_c, the network may decide to discard frames. In this situation it may be advisable to give all DLCIs the following values:

$$B_c + B_e = \text{Access rate} \times T_c$$

Depending on functions implemented on top of the layer 2 core functions, lost frames will be detected and retransmitted, this may be a time-consuming activity

that severely impacts performance. In the latter case, subscribing to high values of CIR, B_c and B_e is important.

5.7.4.2 Communication Rate (CR) and Committed Information Rate (CIR)

The 3746 provides communication rate (CR) support. A part of the access rate (physical line speed) is assigned by the user to each station. IP traffic *per DLCI* is represented by *one* station. The total bandwidth available is split between the stations. This capability differs from the committed information rate (CIR), which is defined as the information rate that the network is committed to transfer under normal conditions over a DLCI.

If all the stations, at any point in time, require more bandwidth than is available, then each station is limited to their user-predefined bandwidth. In case of overflow, the data for those stations that create the overflow are kept in a software queue. They will be transmitted at the next opportunity. If the overflow on the DLCI lasts too long, the data in excess is discarded. The stations that create the overflow are paced and slowed down to their communication rate, while the other ones continue to get their communication rate.

When CR is implemented, if the total physical bandwidth is not fully used the unused bandwidth is available for stations that may then exceed their CR in making use of the unused bandwidth.

For private networks where one wishes to fully utilize all the bandwidth, implementing CR is preferable. For public networks, where CIR is a cost factor and constraint, it is preferable to implement CIR.

The assignment of the communication rate to the stations is done via the CCM at 3746 configuration time. Note that CR and CIR can *not* be shared on the same physical line. Also, line sharing is *not* possible with NCP controlled traffic when CCM has defined the link as having a given CIR. The CIR information is passed by the NNP to the 3746 CLPs at activation time.

When CR is enabled for a given line:

- SNA and APPN FRTEs echo incoming FECNs as BECNs; incoming BECNs reduce XMIT windows as per DYNWIND definitions.
- IP FRTEs ignore incoming FECNs and BECNs.
- FRFHs transport FECNs and BECNs transparently. They do not set BECNs.
- FRTEs and FRFHs set FECNs whenever there is congestion on the transmit physical line.

When CIR is enabled for a given line:

- FRTEs echo incoming FECNs as BECNs.
- FRTEs and FRFHs set FECNs whenever there is congestion on the transmit physical line. COMRATE is still active under CIR so if the physical pipe throughput is reached, the COMRATE process will create BECNs by using B_c as a COMRATE definition. This means that each DLCI's CIR will be reduced so that the sum of all CIRs will equal the access rate, and the relationship between the new CIRs will be the same as the relationship between the B_c s.
- FRFHs transport FECNs and BECNs transparently.
- FRFHs will also set FECNs as soon as the traffic received from a partner subport exceeds that DLCI's CIR during a T_c period. This will also happen when the delay introduced by the FRFH gets too large.

Throughput is optimized when CIR is enabled by tuning to just under the level at which the network sets FECNs and BECNs. All FRTes will adjust their output rate between *minimum information rate* (MIR) and *excess information rate* (EIR). This is called *adaptive-CIR* (A-CIR) and is based on a unique tuning algorithm. This is opposed to CR, which handles the first physical hop and does *not* look for the logical bottleneck within the network. This minimizes queues and delays throughout the network and saves bandwidth on the first hop for other DLCIs that still have end-to-end bandwidth available.

The following provides further detail:

EIR = 0 means the same as $B_c = 0$.

MIR is set at either:

$$\text{MIR} = 0.25 * B_c / T_c, \text{ or}$$

$$\text{MIR} = 9.6 \text{ kbps (if } B_c = 0.)$$

The EIR can be calculated using the following formula:

$$\text{EIR} = (B_c + B_e) / T_c$$

Also,

$$\text{If } B_c = 0, \text{ then } \text{EIR} = B_e / T_c.$$

This is different from the 2216 or 2210, which do not define T_c and therefore commit the whole physical bandwidth as EIR.

FRFHs do not implement A-CIR. They set:

$$\text{CIR} = (B_c + B_e) / T_c$$

If the first hop is congested, they set:

$$\text{CIR} = B_c / T_c$$

Having FRFHs implement CR and FRTes implement CIR will move delays and congestion to the endpoints and optimize the utilization of the network backbone. But it is best that both endpoints implement CIR simultaneously because of the setting of FECNs and BECNs. Note that CIR values need not be the same for these endpoints if there is unbalanced traffic.

Figure 120 on page 182 illustrates how A-CIR is used to maximize the use of available bandwidth between two network endpoints. Every 3.2 seconds a new CIR is computed based on current network congestion (BECNs). (The time interval, 3.2 seconds, was determined pragmatically and took into consideration the anticipated system turnaround time. The value was chosen to exceed it.) This permits learning about current network conditions that resulted from the previous network settings. You will note that this permits the A-CIR curve to stabilize around the available network bandwidth.

You can also see in the graph how the A-CIR curve quickly approaches the network's available bandwidth even after that bandwidth changes. The broken variable CIR curve oscillates and makes measurably less efficient use of the bandwidth.




Figure 120. Adaptive-CIR vs Available Bandwidth

Data points were calculated for the above graph using the respective algorithms for A-CIR and CIR. Available bandwidth was varied to simulate possible network conditions. The interval is 3.2 seconds as mentioned above.

The formula is below:

$$\text{New CIR/Old CIR} = 1 \pm 2^{N-P}$$

where P is PRECISION which is configured and preset to 8. Its possible range is 6 - 10, and N is a run-time variable that is dynamically incremented with a range of 0 - 5. The sign +/- is determined by BECNs. If BECNs are received during the last 1.6 second interval, the sign is - and the CIR is effectively decreased, allowing network congestion to dissipate. If no BECNs are received, the sign is set to +.

When CIR is updated with the same sign as the previous update, N is incremented by 1 until it reaches its maximum value of 5. When the sign changes, N is reset to 0. This is done after computing the new CIR when the CIR is less than the previous CIR and before new CIR computation in the other case.

For the 3746 adapters, either COMRATE or CIR can be specified at the port level. The default value is CIR disabled. This means COMRATE is enabled.

5.7.4.3 Bandwidth Reservation System (BRS)

BRS works in addition to the CIR assigned to a DLCI. It allows the user to define how the CIR of a DLCI should be divided between the different protocols that are using that DLCI. BRS is only used when there is more traffic to send than the available bandwidth, and BRS is not used for FRFH DLCIs.

BRS supports three different traffic types:

- SNA (APPN/DLUR, and HPR-ERP) with link-level error recovery
- HPR-ERP (without link-level error recovery)
- IP

IP traffic can further be differentiated in that five IP sockets can also be defined. The total bandwidth assigned may not exceed 100% of the CIR.




Figure 121. 3746 Frame Relay BRS

When CIR is enabled, T_c , B_c , and B_e can be defined for each DLCI. B_c , or B_e may be equal to zero, but B_c plus B_e must be greater than zero. B_c and B_e are expressed in multiples of DATABLK.

For FRFH and FRTE DLCIs:

T_c The measurement interval (T_c) is defined at the DLCI level. The default value is 0.1 second. The value is specified in tenths of a second (1-255).

B_c Committed burst size in units of DATABLK (0-64).

B_e Excess burst size in units of DATABLK (0-64).

Then for FRTE DLCIs, each protocol may be assigned a percentage of those CIR values:

```

BRS
DLUR/APPN/HPR (ERP):      T%
HPR (non-ERP):            U%
IP:                       V%
IP socket #1:             W%
.
.
IP socket #5:             Z%

```

The percentages can be defined between 0% and 100%. A zero value means that no bandwidth is reserved for that protocol. This means that all this traffic will be discarded when congestion occurs. If a protocol has no BRS defined, then that protocol's traffic does not participate in BRS; all that protocol's traffic is transmitted whatever the level of congestion.

5.7.5 Consolidated Link Layer Management (CLLM)

CLLM is an optional FR management function that is not widely supported by the industry but it has been adopted by some frame relay switch manufacturers. CLLM provides some of the same management information provided by LMI, in particular, outage notification. CLLM's main use is to provide asynchronous congestion notification to attaching devices. A single CLLM message may indicate outage or congestion for multiple PVCs. The frame relay protocol supports the following standards for CLLM:

1. ANSI T1.618
2. ITU-T (CCITT) Q.922 Annex A
3. ITU-T (CCITT) X.36 Annex C

On the 3746 from microcode level D46133 onwards, CLLM messages are received and acted on by the 3746.

5.7.5.1 Using CLLM

If your network provider supports CLLM, you can configure 3746 Frame Relay to throttle down its transmit rate for PVCs contained in a CLLM message. CLLM messages contain a cause code that indicates the type and severity of the problem being reported. The device reacts differently depending on the cause code and the CIR configured for each PVC contained in the CLLM message. When the device receives a CLLM message that indicates:

- A short-term condition, and the configured CIR for the PVC is nonzero, the frame relay protocol will throttle the transmit rate for the affected PVCs by the configured information rate (IR) decrement percentage.
- A long-term condition, the frame relay protocol will set the transmit rate for the affected PVCs to the calculated minimum IR.
- Facility or equipment failure or maintenance action, or if the CIR was configured as zero, the FR protocol will continue to transmit any queued data for the affected PVCs but will not accept any more outgoing packets from the upper layer protocols until the congestion condition is cleared.

Once a CLLM message for a PVC has been received, if the device does not receive any CLLM messages or BECNs within the Ty timer period or if a frame

without a BECN is received, the device will consider the congestion condition cleared and gradually return the PVC to its configured transmission rates.

If you are using CLLM to control congestion, you must not configure DLCI 1007 for any other use.

5.7.5.2 Configuring CLLM

CLLM support can be configured from CCM and the CONFIG line command interface.

5.7.6 Frame Relay Port Sharing

Frame relay is the only type of serial interface that is supported by all 3745 and 3746 protocol stacks. A single frame relay interface (port) can be activated simultaneously by each protocol stack, 3746 NN/DLUR and IP and 3745 NCP. Frame relay together with the 3745 and 3746, is ideally suited for building a multiprotocol network backbone.

Notes:

1. When using a dual-CCU 3746-900, a frame relay line can only be activated from either CCU-A (NCPA) or CCU-B (NCPB), not both at the same time.
2. NCP 3746-900 frame relay support is limited to SNA only. NCP does not support IP over frame relay on 3746-900 attachments.

5.7.7 Frame Relay PVC Sharing

To accomplish IP, APPN, and NCP/subarea connectivity one can use either separate PVCs for each of the individual protocol stacks or use the 3746-9x0 PVC sharing facilities.

The 3746-9x0 PVC sharing facilities enable the 3746-9x0 to distinguish between and share a PVC for 3746 IP, 3746 NN/DLUR (routed and bridged frame format), and NCP/Boundary (routed and bridged frame format) traffic. The sharing of a 3746 controlled frame relay line with NCP requires NCP V7R5.

Internal routing of incoming frames is done as follows:

1. Check DCLI number. (This may be an FRFH DLCI.)
2. Is the traffic IP (NLPID=X'CC')?
3. Routed frame format (NLPID=X'08'):
 - L2 field indicates ERP or non-ERP.
 - L3 field indicates INN, BNN or APPN, or HPR.
4. Bridged frame format (NLPID=X'80').

INN traffic always goes to the NCP; BNN/APPN traffic is routed by the DSAP. HPR traffic is routed by ANR label, therefore HPR packets can be switched at the adapter level in the 3746 without being routed to software components.

5.7.7.1 Frame Relay Frame Handler Functions

The NCP frame relay frame handler functions (FRFH) take the traffic arriving on a VC (3745 or 3746 port), and switch all the traffic on that VC to an outbound VC. All forms of encapsulation are supported as the frame contents (apart from the frame header) are not examined by the FRFH function.

The NCP FRFH function can switch VCs between ports on the 3745 and ports on the 3746.

The 3746 frame relay frame handler function is controlled and defined from CCM, and can switch frame relay traffic between VCs on 3746 adapters.




Figure 122. 3745 Frame Handler Function

5.7.7.2 IP over Frame Relay

3745 and 3746 IP traffic is sent using NLPID encapsulation with NLPID=X'CC', and uses the RFC1490 routed frame format. For IP the bridged frame format is not supported. The IP SAP X'AA' is used for LAN traffic but is not used for IP over frame relay. IP traffic from 3746 adapters is passed to the 3746 IP component; traffic from the 3745 adapters is passed to the 3745 IP component.

IP traffic on 3746 adapters can be multiplexed with APPN and BNN traffic on the same DLCI. SNA traffic is encapsulated with NLPID=X'08' (routed frame format) or NLPID=X'80' (bridged frame format). Traffic for the 3746 APPN CP and NCP BNN function is distinguished by the SAP values in incoming SNA frames.

Figure 123 on page 187 shows the IP support.




Figure 123. 3745 and 3746 IP over Frame Relay

5.7.7.3 INN over Frame Relay

Frame relay INN traffic can be either in the RFC1490 routed or bridged frame format. INN traffic on 3745 ports can share a DLCI with IP traffic for the NCP. This traffic can be distinguished by the NLPIDs used (see Figure 124 on page 188). INN traffic over frame relay (routed or bridged frame format) always uses DSAP=X'04' and SSAP=X'04'. Traffic from 3746 ports can also be passed to the NCP INN function. A single INN station per DLCI is supported whether on 3745 or 3746 adapters.




Figure 124. 3745 INN Traffic - Routed Frame Format




Figure 125. 3745 INN Traffic - Bridged Frame Format

5.7.7.4 BNN and APPN over Frame Relay

BNN and APPN traffic from 3745 adapters uses a separate DLCI and cannot be multiplexed with other traffic. BNN and APPN traffic from 3746 adapters can be multiplexed with IP traffic on the same DLCI (see Figure 126 and Figure 127).




Figure 126. BNN/APPN Traffic - Routed Frame Format




Figure 127. BNN/APPN Traffic - Bridged Frame Format

5.7.8 Frame Relay - General Configuration

To define a frame relay port, you must specify what type of traffic (APPN, and/or IP, and/or FRFH) will be transported, and IP and APPN names for the port (see Figure 128).

Port Configuration

Configure a Port

DLC type: ☒ Frame Relay ☐ PPP ☐ SDLC ☐ X.25

Network: ☒ APPN ☒ IP ☐ FRFH Port: 2443

APPN name: FR2443 IP name: IP2443

Comments (optional):

Ports Already Configured

Port	APPN name	IP name	DLC type	No. of stations
2445	FR2445	FR	0	
2446	FR2446	FR	0	
2444	IP2444	FR	0	
2443	FR2443	IP2443	FR	0

Buttons: Add, Modify, Copy..., Delete, DLC parameters..., APPN parameters..., DLCI..., APPN/IP Stations, IP parameters..., OK, Search..., Search next, Cancel, Help

Figure 128. Frame Relay Port Configuration

For each frame relay port, you must also configure DLC parameters. You may specify the type of congestion control, COMRATE or CIR, used for frame relay on this panel too. To use bandwidth reservation on this port, CIR must be selected. COMRATE is mandatory if the port will be shared with NCP.

Frame-Relay Port - DLC Parameters

Port: 2444 Name:

Port Type: ☒ Leased ☐ Switched

Interface: ☒ V.24 ☐ V.25B ☐ V.35 ☐ X.21

Clocking: ☐ Internal ☐ Direct ☒ External

Data Rate: ☐ High ☒ Low

Speed: 9.6 kbps

Transmit NRZI? ☐ Yes ☒ No

Interframe gap (ADDIFG)? ☐ Yes ☒ No

Bandwidth management: ☒ COMRATE ☐ CIR

Maximum frame size (MAXFRAME): 2106 bytes [282-8050]

Data block size (DATABLK): 2048 bytes [265-16732]

Enable timer (ENABLT0): 22 tenths seconds [1-16320]

Disable timer (DSABLT0): 30 tenths seconds [1-16320]

Boundary node identifier: 4FFF00000000 hexadecimal [0-7FFFFFFFFF]

Local SAP (LSAP): 8 hexadecimal [04-9C]

Buttons: OK, LMI..., CIR parameters..., Save as defaults, Cancel, Help

Figure 129. Frame Relay DLC Configuration

5.7.9 Frame Relay - IP Configuration

There are two methods available for defining IP over a frame relay DLCI. The first method is to define all parameters; the second is to allow incoming calls from the network, in which case all definitions are then made dynamically.

5.7.9.1 Static Definitions

On each defined frame relay port, multiple PVCs may be defined, each PVC is identified by its *data link control identifier* (DLCI) number. For each DLCI transporting IP, an IP name is needed, and a remote IP address may be necessary. A remote IP address must be defined if the parameters *Enable multicast emulation* and *Enable protocol broadcast* in the dialog *IP over frame-relay parameters* have not been selected. Pre-defining the remote IP address allows the 3746 to route data to a remote IP address over the correct DLCI, although no dynamic routing protocol has supplied the route to that address (see “Defining the IP Address” on page 192). Otherwise, the data is routed over each DLCI supporting IP on that port.

When running IP over frame relay, the DLCI number is the hardware address, much the same as the MAC address of a token ring station.

For each APPN station, and for IP, the COMRATE value can be specified on this panel.

For each frame relay port, IP over frame relay parameters must be defined if IP will be routed over this port. Here either a single local IP address or multiple local IP addresses associated with this port are defined.

Next hop awareness may also be enabled. This allows the 3746 IP router to obtain information from the CLP about which remote IP addresses are reachable over which DLCI.

IP Over Frame-Relay Parameters

Port: 2444 Name: IP2444

General

Automatic reactivation? ☒ Yes ☐ No

☒ Enable orphan circuits ☒ Enable multicast emulation

☒ Enable protocol broadcast ☒ Enable congestion monitor

Down time: 0 seconds [0-300]

Addresses

☐ Enable next hop awareness (NHA)

IP address: 9.24.105.1 Subnet mask: 255.255.255.0

Add Modify Delete

NHA	9.24.104.1	255.255.255.0
	9.24.105.1	255.255.255.0

OK Save as defaults Cancel Help

Figure 130. Frame Relay IP Parameters

Port: 2444 Name: IP2444

Configure a DLCI

Network: ☐ APPN ☒ IP ☐ FRFH DLCI number: 32 num. [16-991]

DLCI IP Name: IP244401 Remote IP address:

Communication rate (COMRATE): 16384 bits [16384-1048576]

DLCIs Already Configured

Network	DLCI no.	APPN COMRATE	IP COMRATE
IP	32	N/A	16384

Buttons: Add, Modify, Delete, APPN stations..., FRF set..., OK, Default DLCI..., Cancel, Help

Figure 131. Frame Relay IP DLCI

Defining the IP Address: On each port one or more local IP addresses must be defined. The IP addresses are defined per port not per DLCI (see Figure 130 on page 191).

Up to now we have an active PVC identified by its DLCI and we have defined our own IP address. But we don't know the IP address of the other end of the PVC. However we know the hardware address (DLCI).

How do we get the IP address at the other end of the PVC? This problem can be solved in one of two ways:

- Static definition. This means you need to assign the remote IP address at the end of each DLCI (see Figure 131).
- Dynamic address resolution.

This means the IP protocol stack itself gets the IP address from the remote station by using *inverse ARP* (IARP) on a PVC. IARP is specially adapted to the requirements of frame relay (see RFC 1490 for details). For ARP to be supported over frame relay links, *multicast emulation* must be enabled. When enabled, the 3746 can emulate multicast frames by sending the frame over each DLCI supporting IP on a specific port.

See Figure 244 on page 279 for how to dump the IARP table and see Figure 245 on page 279 for how to display all user-defined address pairs.

Protocol Broadcast: If we want to use protocol broadcasts such as PING 10.255.255.255, it is necessary to enable the protocol broadcast feature as well (see Figure 130 on page 191). This allows the 3746 to emulate a broadcast by sending packets down each DLCI on a port, thereby reaching all partners.

5.7.9.2 Dynamic Definitions

To support dynamic definitions, enable Orphan Circuit. This allows the port to accept any PVC establish request from the frame relay net (Figure 130 on page 191).

5.8 Point-to-Point Protocol (PPP)

An alternative to using frame relay to transport IP traffic over 3746 serial line connections is to use PPP. RFC 1331 - *The Point-to-Point Protocol (PPP)* describes the transmission of multiprotocol datagrams over point-to-point links. PPP was derived from earlier work on SLIP (Serial Line IP, see RFC 1052 - *Transmission of IP Datagrams over Serial Lines*). PPP is a more robust, full-featured protocol that solves many of the deficiencies within SLIP.

Learning from the shortcomings of SLIP, the developers started with the concept that all standards need to be extensible. Not all vendors need the same features in every product, and new features can be added easily over time without affecting older implementations.

On PPP, the communication of implementation's features set is done through negotiation. Each end of the link advertises which features it has implemented, and the peers agree on a common set of features to be used. This allows interoperability among a wide range of products. The most important aspect of PPP negotiation is that it is automatic, which can result in considerable savings of human configuration time.

PPP is suited to encapsulate protocols other than just TCP/IP. However, the 3746 IP router uses PPP for IP transport only. PPP automatically discovers which network protocols are supported through negotiation.

HDLC is the default framing described in RFC 1331 that enables the 3746 IP to use the 3746-9x0's existing bit-synchronous hardware.

PPP is comprised of three main components:

- A method for encapsulating datagrams over serial links

The 3746 IP router supports RFC 1332 - *PPP Internet Protocol Control Protocol (IPCP)* for the encapsulation of IP over PPP connections.

- Link Control Protocols (LCPs)

The LCP is used for establishing, configuring, and testing data-link connections.

- Network Control Protocols (NCPs)

Network Control Protocols (NCPs) are protocols used for establishing and configuring different network layer protocols.

To establish communications over a point-to-point link, each end of the PPP link first sends LCP packets to configure and test the data link. After the link has been established, the peer may be authenticated (see 5.8.2, "Link Control Protocol (LCP)" on page 194). Then PPP sends NCP packets to choose and configure one or more network layer protocol. When each chosen network layer protocol has been configured, data can be sent on the link.

5.8.1 Internet Protocol Control Protocol (IPCP)

PPP IPCP options on the 3746 IP router are:

- Compression support

The 3746 IP implementation supports Van Jacobson header compression. IP header compression allows PPP to scale for slower links. Enabling compression (optional) should be considered on low-speed connections.

- Sending local IP address

A configuration option exists to send the local IP address to the remote end of the link. This option should be enabled if the other end requires the IP address.

- Request remote IP address

A configuration option exists to request the IP address of the remote end of the link. If the remote IP address is received, it will be displayed in PPP monitoring statistics.

5.8.2 Link Control Protocol (LCP)

The LCP is used to agree automatically upon the encapsulation format options, handle varying limits on packet sizes, decide authentication and compression techniques, determine when a link is functioning properly and when it is down, detect a looped-back link, and terminate the link.

PPP LCP options on the 3746 IP router are:

- Maximum receive unit (MTU)

This sets the maximum packet size of the information field transferred in a single datagram.

- Magic number

This specifies whether or not the magic number option is enabled. The magic number option provides a way of detecting looped-back links.

The magic number is a random number added to LCP configure requests. When the LCP receives a configure request with a magic number, this number is compared with the last magic number sent. If the two magic numbers are different, the link is not considered looped back. If the magic numbers are the same, the link is brought down and up again.

- Configure Request retry timer
- Configure Request retry count
- Not Acknowledged retry count
- Terminate Request retry count

RFC 1331 specifies the use of Password Authentication Protocol (PAP) or Challenge-Handshake Authentication Protocol (CHAP) to enforce link security. However, neither authentication option has been implemented on the 3746 IP router.

5.8.3 Network Control Protocol (NCP)

PPP NCP configurable options on the 3746 IP router are:

- Configure Request retry timer
- Configure Request retry count
- Not Acknowledged retry count
- Terminate Request retry count

Note: LCP and NCP use the same negotiation technique and similar timer and counters apply. An LCP connection must become active before an NCP connection can be started.

5.8.4 PPP BRS on the 3746 Base Adapters

3746 BRS on PPP connections work as a subset of the 3746 multiaccess enclosure BRS. The bandwidth of a PPP connection can be divided up between five pre-defined traffic types:

- IP
- Rlogin
- Telnet
- SNMP
- Multicast

The bandwidth of the PPP connection can be assigned as a percentage to user-defined traffic classes. Then each protocol type can be assigned to one of the traffic classes.

5.8.4.1 Priority Levels

Within each traffic class, the different protocols can also be assigned a *priority level*. The following priority levels are defined:

- Urgent (U)
- High (H)
- Normal (N)
- Low (L)

Figure 132 on page 196 shows three traffic classes, each traffic class has its own set of data which has been given a priority shown by the four queues, urgent(U), high (H), normal (N), and low (L).




Figure 132. Traffic Class Queues

Figure 132 shows three traffic classes: class A with 50% of the DLCI CIR, class B with 10% of the DLCI CIR, and class C with 40% of the DLCI CIR. Traffic bound for the link shown is differentiated by the previously discussed protocol types or filters, and is assigned one of the four priorities.

5.8.5 PPP Basic Configuration Procedure

This section outlines the minimum configuration steps that you are required to perform to get the PPP protocol up and running.

Defining PPP connections includes the following:

- Add the PPP interface
- Configure LCP options
- Configure NCP options
- Configure Bandwidth Reservation over PPP

In addition to the basic PPP configuration, which must be done, you can also configure BRS (an optional feature) over PPP. For information on bandwidth reservation, see 5.7.4.3, “Bandwidth Reservation System (BRS)” on page 182.

5.8.6 PPP - General Configuration

To define a serial port for PPP, the port number and IP name must be specified. This port cannot be shared with protocols other than IP.

Port Configuration

Configure a Port

DLC type: ☐ Frame Relay ☒ PPP ☐ SDLC ☐ X.25

Network: ☐ APPN ☒ IP ☐ FRFH Port: 3024

APPN name: IP name:

Comments (optional):

Ports Already Configured

Port	APPN name	IP name	DLC type	No. of stations
3024	PPP3024	PPP		

Figure 133. PPP Port Configuration

5.8.7 PPP - IP Configuration

The PPP port parameters 1/2 display (Figure 135 on page 198) shows how to configure the physical port for PPP, set the MTU size, and how to specify the port IP address. The *unnumbered IP address* check box is also found here, this is a facility to save IP addresses. Normally a serial link is configured by using two IP addresses, one on each end of the link. These two addresses reserve a whole IP address space. That is, if you use a class C network address, 252 IP addresses would be reserved (the whole net) although only two addresses are used. If you want to use unnumbered addresses, be sure that both machines on this link support unnumbered addressing (see Figure 134 on page 198). Also, if no IP addresses are assigned then the link cannot be managed by IP management software.




Figure 134. PPP Unnumbered IP Address

The screenshot shows the 'PPP Port - Parameters 1/2' dialog box. The 'Port' is 3024 and the 'Name' is IP3024. The 'Port Type' is 'Leased'. The 'Interface' is 'V.24'. The 'Clocking' is 'Internal'. The 'Data Rate' is 'High'. The 'Speed' is '9.6 Kbps'. The 'Transmit NRZI?' option is 'No'. The 'Interframe gap (ADDIFG)?' option is 'No'. The 'Automatic reactivation?' option is 'Yes'. The 'Maximum transmission unit' is '2048 bytes [576-4088]'. The 'Port IP Address' section shows 'Unnumbered IP address' selected, with 'IP address' set to '9.24.104.1' and 'Subnet mask' set to '255.255.255.0'. The 'OK' button is highlighted.

Figure 135. PPP DLC Parameters 1/2

LCP Parameters can be defined on the PPP DLC parameters 2/2 dialog shown in Figure 136 on page 199.

Port: 3024 Name: IP3024

LCP Parameters

Retry timer: 3000 milliseconds [200-30000]

Config tries: 20 numerical [1-100]

NAK tries: 10 numerical [1-100]

Terminate tries: 10 numerical [1-100]

Down time: 0 seconds [0-300]

Enable timer (ENABLT0): 22 ts [1-16320]

☒ Magic number

☐ IP compression Number of slots: 16 num [1-16]

☐ Send IP address ☐ Request IP address

OK Save as defaults Cancel Help

Figure 136. PPP DLC Parameters 2/2

To configure network control protocol parameters, select IP parameters from the Configuration menu in CCM, then select **PPP NCP Parameters**. Figure 137 shows the NCP parameters dialog.

PPP NCP Default Parameters

Retry timer: 3000 milliseconds [200-30000]

Config tries: 20 numerical [1-100]

NAK tries: 10 numerical [1-100]

Terminate tries: 10 numerical [1-100]

OK Cancel Help

Figure 137. PPP NCP Parameters

5.9 X.25 Protocol

The 3746-9x0 running X.25 protocol can be either:

- A DTE attached to a packet switched data network (PSDN) that conforms to ITU-T Recommendation X.25 (1993).
- A DTE directly attached to another DTE that conforms to the ISO 7776 (layer 2) and ISO 8208 (layer 3) standards.

For more detailed information on 3746 X.25 support, see *3745 Communication Controller Models A, 3746 Multiprotocol Controller Models 900 and 950 Planning Guide*, GA33-0457.

The X.25 ITU-T recommendation specifies the following three OSI layers:

- Layer 1: Physical layer
- Layer 3: Data link layer, called Link Access Procedure Balanced (LAP-B)
- Layer 3: Network layer, called Packet Layer Protocol (PLP)

For X.25, the physical layer is the same as for any other WAN protocol implemented in the CLP (frame relay, SDLC, or PPP). The X.25 support implements the X.25 layer 2 (LAP-B) and 3 (PLP). It also implements another layer above PLP for adapting SNA, APPN/HPR, and IP protocols to the X.25 layers. Figure 138 on page 200 shows how the X.25 functions are implemented in the 3746-9x0. The CLP can be under NCP or NNP control.




Figure 138. X.25 Processing Layers

The following TCP/IP routing protocols are supported over X.25:

- RIP
- OSPF
- BGP

Neighbors around the X.25 network must be configured manually through CCM.

The interface layer between X.25 and the IP control point (IP-PLP) conforms to RFC 1356.

5.9.1 X.25 Port Sharing

X.25 ports can be shared by APPN and IP traffic, and a single NCP (X.25 ODLC). In the case of a dual CCU 3745, only a single NCP can activate a 3746 X.25 link at any one time. NPSI traffic cannot share an X.25 line with other protocols.

NCP V7R4 introduced support for transporting SNA traffic over X.25 without requiring NPSI.




Figure 139. X.25 Port Sharing

5.9.2 X.25 - General Configuration

To enable a serial port for X.25, the protocols (either APPN or IP) that will use the port, and port names must be specified.

Port Configuration

Configure a Port

DLC type: ☐ Frame Relay ☐ PPP ☐ SDLC ☒ X.25

Network: ☒ APPN ☒ IP ☐ FRFH Port: 3024

APPN name: APPN3024 IP name: IP3024

Comments (optional):

Ports Already Configured

Port	APPN name	IP name	DLC type	No. of stations
3024	APPN3024	IP3024	X25	0

Buttons: Add, Modify, Copy..., Delete, DLC parameters..., APPN parameters..., DLCL..., APPN/IP Stations, IP parameters..., OK, Search..., Search next, Cancel, Help

Figure 140. X.25 Port Configuration

5.9.3 X.25 - IP Configuration

The local IP addresses for an X.25 port are defined as shown in Figure 141. Multiple IP addresses may be defined on a single port, but each must belong to a separate subnet.

IP address:	Subnet mask:
9.24.108.1	255.255.255.0
9.24.105.1	255.255.255.0
9.24.108.1	255.255.255.0

Figure 141. X.25 Port IP Configuration

5.9.4 X.25 - IP Configuration Example

Figure 142 on page 203 and the configuration screens following that show an X.25 configuration involving two 3746 machines, and how to create definitions for this configuration using CCM. The machines are connected by direct DTE-to-DTE connections; there is no X.25 network between the machines. It gives an example of how to configure IP, and shows how IP shares the port with APPN. Two virtual circuits are defined: one PVC dedicated to APPN, and one SVC dedicated to IP.

In the DTE-to-DTE environment, the logical channel numbers (LCNs) must match at both ends of the physical line. When connected to an X.25 network, the LCNs must match those subscribed to in the network.

Both port IP addresses must be in the same subnet, and subnet masks must match.

In the DTE-to-DTE environment, the calling DTE address must be included in the call request. This is specified in Figure 153 on page 209 and Figure 164 on page 214. Also note that the Type of Address (TOA) was set to Network Dependent (0) and that Numbering Plan Identification (NPI) was set to X.121 (3). The value 03 can be seen prepended to the DTE address in Figure 150 on page 207.




Figure 142. X.25 Example Configuration

5.9.4.1 X.25 IP Example - 3746 B

Follow these steps to configure IP over X.25:

1. Configure port 2272 for IP (and APPN if desired).

Port Configuration

Configure a Port

DLC type: ☐ Frame Relay ☐ PPP ☐ SDLC ☒ X.25

Network: ☒ APPN ☒ IP ☐ FDDI Port: 2272

APPN name: X252272A IP name: X252272I

Comments (optional):

Buttons: Add, Modify, Copy...

Ports Already Configured

Port	APPN name	IP name	DLC type	No. of stations
2272	X252272A	X252272I	X25	3

Buttons: Delete, DLC parameters..., APPN parameters..., DLCI..., APPN/IP Stations, IP parameters...

Buttons: OK, Search..., Search next, Cancel, Help

Figure 143. X.25 Port Configuration

2. Select **DLC parameters**.
3. In this case as the machines are directly connected, select **Direct Clocking**. The other machine must be configured for **External Clocking**.

X.25 Port - DLC Parameters

Port: 2272 Name: APPN: X252272A IP: X252272I

Port Type: ☒ Leased ☐ Switched

Interface: ☐ V.24 ☐ V.25B ☐ V.35 ☒ X.21

Clocking: ☐ Internal ☒ Direct ☐ External

Data Rate: ☒ High ☐ Low

Speed: 2048 kbps

Transmit NRZI? ☐ Yes ☒ No

Interframe gap (ADDIFG)? ☐ Yes ☒ No

Enable timer (ENABLT0): 3000 tenths seconds [1-16320]


Disable timer (DSABLT0): 3000 tenths seconds [1-16320]

Buttons: LAPB parameters..., PLP parameters...

Buttons: OK, Save as defaults, Cancel, Help

Figure 144. X.25 Port DLC Parameters

4. Select **LAPB Parameters**.
5. Set Role=DTE. The other machine must be configured for Role=DCE.
6. Select **OK**.



X.25 Port - LAPB Parameters

Port: 2272 Name: X252272A

DTE to DTE? ☒ Yes ☐ No LAPB role: DTE Modulo: ☒ 8 ☐ 128

K - Frame window size (MWINDOW): 7 numerical [1-7]

N1 - Maximum frame size (FRMLGTH): 131 bytes [35-4100]

T1 - Reply timer (TPTIMER): 50 tenths seconds [5-255]


T2 - Acknowledgement timer: 5 tenths seconds [0-255]

T4 - Inactivity timer: 60 seconds [1-65535]

N2 - Maximum transmission attempts (NPRETRY): 7 numerical [1-255]

Figure 145. X.25 LAPB Parameters

7. Select **PLP parameters**.
8. Set Local DTE Address=155522721234567.
9. Set PLP Role=Negotiable.
10. Set Logical Channel Numbers as shown.



X.25 Port - PLP Parameters 1/2

Port: 2272 Name: X252272A

X.25 Local DTE Address

X.25 local DTE address: 155522721234567

General PLP Parameters

PLP role: Negotiable X.25 version: ITU-T 93


Accounting required: ☐ Yes ☒ No Segment tax unit: 64

Logical Channel Numbers (LCN) [1-4095]

Permanent Virtual Circuits (PVCs)		DTE Incoming Channels (SVCs)	
Lowest: 1	Highest: 1	Lowest:	Highest:
Two-Way Channels (SVCs)		DTE Outgoing Channels (SVCs)	
Lowest: 2	Highest: 2	Lowest:	Highest:

Figure 146. X.25 PLP Parameters 1/2

11. Select **PLP parameters 2/2**.
12. Set PLP Modulo=8.
13. Select **OK** three times. This brings you back to the screen shown in Figure 143 on page 204.



X.25 Port - PLP Parameters 2/2

Port: 2272 Name: X252272A

PLP modulo: ☒ 8 ☐ 128

SVC Defaults

Packet window calling to called: 2 numerical [1-7]

Packet window called to calling: 2 numerical [1-7]

Packet size calling to called: 128 bytes

Packet size called to calling: 128 bytes

PLP Port Timers (seconds [1-255])

Restart timer (T20): 180 Reset timer (T22): 180

Call timer (T21): 200 Clear timer (T23): 180

PLP Port Maximum Retries (numerical [0-255])


Restart retry (R20): 1 Clear retry (R23): 1

Reset retry (R22): 1

OK Save as defaults Cancel Help

Figure 147. X.25 PLP Parameters 2/2

14. Select **IP Parameters**.
15. Set MTU as desired.
16. Set IP Address=145.17.4.61. (This is the local IP address.)
17. Set IP Mask=255.255.255.0.
18. Select **OK**.



IP Over X.25 - Port Parameters

Port: 2272 Name: X252272I

General

Automatic reactivation? ☒ Yes ☐ No

Maximum transmission unit: 2058 bytes [576-4096]

☐ No idle timer Idle timer: 30 seconds [1-255]

Addresses

IP address:	Subnet mask:
145.17.4.61	255.255.255.0

Buttons: Add, Modify, Delete

OK Save as defaults Cancel Help

Figure 148. X.25 Port Parameters

19. Select **APPN/IP Stations**.
20. Create an APPN station if desired.
21. Select **ADD**.

X.25 Station Configuration

Port: 2272 Name: X252272A

Configure an X.25 Station

Network: ☒ APPN ☐ IP Station name: AP2272P ☒ PVC ☐ SVC **Add**

TOA/NPI? ☒ Yes ☐ No TOA: Network dependent (0) NPI: X.121 (3) **Modify**

Remote DTE address: LCN: 1 PU type: ☐ 2 ☒ 2.1 **Copy...**

Comments (optional):

X.25 Stations Already Configured

Network	Station name	Type	DTE address or LCN
APPN	AP2272P	PVC 2.1	1

Delete **Search...**

Search next

DLC parameters...

APPN parameters...

IP parameters...

User facilities...

OK **Save as defaults** **Cancel** **Help**

Figure 149. X.25 APPN Station Configuration

22. Create an IP station.
23. Set IP Station Name=IP2272S.
24. Set SVC.
25. Set TOA/NPI=YES.
26. Set Remote DTE address=144423041234567.
27. Select **Add**.
28. Select **OK**.

X.25 Station Configuration

Port: 2272 Name: X252272I

Configure an X.25 Station

Network: ☐ APPN ☒ IP Station name: IP2272S ☐ PVC ☒ SVC **Add**

TOA/NPI? ☒ Yes ☐ No TOA: Network dependent (0) NPI: X.121 (3) **Modify**

Remote DTE address: 144423041234567 LCN: 1 PU type: ☐ 2 ☒ 2.1 **Copy...**

Comments (optional):

X.25 Stations Already Configured

Network	Station name	Type	DTE address or LCN
APPN	AP2272P	PVC 2.1	1
IP	IP2272S	SVC N/A 03	144423041234567

Delete **Search...**

Search next

DLC parameters...

APPN parameters...

IP parameters...

User facilities...

OK **Save as defaults** **Cancel** **Help**

Figure 150. X.25 IP Station Configuration

29. Select **IP parameters** for station IP2272S.
30. Set Remote IP Address=145.17.4.71.
31. Select **OK**.

IP Over X.25 - Station Parameters

Port: 2272 Name: X252272I
 Station name: IP2272S
 X.25 remote DTE address: 03 144423041234567
 PVC logical channel number: N/A

General

Automatic reactivation? ☐ Yes ☒ No

Addresses

Remote IP address:

Figure 151. X.25 IP Station Parameters

32. Select **User Facilities**.

33. Set Add Calling DTE Address.

X.25 Station - SVC Call Requests User Facilities and Data

Port: 2272 Name: X252272I Network: IP
 Station name: IP2272S X.25 remote DTE address: 03 144423041234567

Protocol identifier: ☐ 'CB' ☐ 'C3' ☒ 'CC'

☐ Window size negotiation facility In: 2 Out: 2 num [1..7]
☐ Packet size negotiation facility In: 128 Out: 128 bytes
☐ Request reverse charging facility
☐ Request inter-network call redirection and deflection (ICRD) control facilities
 ICRD status selection: ☒ Prevention requested ☐ Allowance requested
☒ Add calling DTE address
☐ Throughput class negotiation facility
 Format: ☒ Basic ☐ Extended In: 192 Out: 192 Kbps
☐ Closed user group selection facility ☐ With outgoing access selection
 Format: ☒ Basic ☐ Extended 1 numerical [0..99]

Figure 152. X.25 User Facilities

34. Select **Calling DTE address parameters**.

35. Set Calling DTE address=155522721234567.

36. Select **OK**.

X.25 Station - SVC Calling DTE Address

Port: 2272 Name: X252272I Network: IP
 Station name: IP2272S
 X.25 remote DTE address: 03 144423041234567

Calling DTE Address—

TOA: Network dependent (0) NPI: X.121 (3)
 Calling DTE address: 155522721234567

OK Cancel Help

Figure 153. X.25 Calling DTE Address Parameters

5.9.4.2 X.25 IP Example - 3746 A

Follow these steps to configure IP over X.25:

1. Configure port 2304 for IP (and APPN if desired).

Port Configuration

Configure a Port—

DLC type: ☐ Frame Relay ☐ PPP ☐ SDLC ☒ X.25

Network: ☒ APPN ☒ IP ☐ FRN Port: 2304

APPN name: X252304A IP name: X252304I

Comments (optional): X25 Line to ERS5 IP X25

Buttons: Add, Modify, Copy...

Ports Already Configured—


Port	APPN name	IP name	DLC type	No. of stations
2304	X252304A	X252304I	X25	2

Buttons: Delete, DLC parameters..., APPN parameters..., DLC..., APPN/IP Stations, IP parameters...

OK Search... Search next Cancel Help

Figure 154. X.25 Port Configuration

2. Select **DLC parameters**.
3. In this case as the machines are directly connected, select **External Clocking**. The other machine must be configured for Direct Clocking.



X.25 Port - DLC Parameters

Port: 2304 Name: APPN: X252304A IP: X252304I

Port Type
☒ Leased
☐ Switched

Interface
☐ V.24
☐ V.25B
☐ V.35
☒ X.21

Clocking
☐ Internal
☐ Direct
☒ External

Data Rate
☒ High
☐ Low

Speed: 2048 kbps
 Transmit NRZI? ☐ Yes ☒ No
 Interframe gap (ADDIFG)? ☐ Yes ☒ No


Enable timer (ENABLT0): 3001 tenths seconds [1-16320]
 Disable timer (DSABLT0): 3000 tenths seconds [1-16320]

LAPB parameters... PLP parameters...

OK Save as defaults Cancel Help

Figure 155. X.25 Port DLC Parameters

4. Select **LAPB Parameters**.
5. Set Role=DCE. The other machine must be configured for Role=DTE.
6. Select **OK**.



X.25 Port - LAPB Parameters

Port: 2304 Name: X252304A


DTE to DTE? ☒ Yes ☐ No LAPB role: DCE Modulo: ☒ 8 ☐ 128

K - Frame window size (MWIND0W): 7 numerical [1-7]
 N1 - Maximum frame size (FRMLGTH): 131 bytes [35-4100]
 T1 - Reply timer (TPTIMER): 50 tenths seconds [5-255]
 T2 - Acknowledgement timer: 4 tenths seconds [0-255]
 T4 - Inactivity timer: 60 seconds [1-65535]
 N2 - Maximum transmission attempts (NPRETRY): 7 numerical [1-255]

OK Save as defaults Cancel Help

Figure 156. X.25 LAPB Parameters

7. Select **PLP parameters**.
8. Set Local DTE Address=144423041234567.
9. Set PLP Role=Negotiable.
10. Set Logical Channel Numbers as shown.



X.25 Port - PLP Parameters 1/2

Port: 2304 Name: X252304A

X.25 Local DTE Address
X.25 local DTE address: 144423041234567

General PLP Parameters

PLP role: **Negotiable** X.25 version: **ITU-T 93**

Accounting required: ☐ Yes ☒ No Segment tax unit: 64

Logical Channel Numbers (LCN) [1-4095]

Permanent Virtual Circuits (PVCs)
Lowest: 1 Highest: 1

DTE Incoming Channels (SVCs)
Lowest: Highest:


Two-Way Channels (SVCs)
Lowest: 2 Highest: 2

DTE Outgoing Channels (SVCs)
Lowest: Highest:

OK PLP parameters 2/2... Save as defaults Cancel Help

Figure 157. X.25 PLP Parameters 1/2

11. Select **PLP parameters 2/2**.
12. Set PLP Modulo=8.
13. Select **OK** three times. This brings you back to the screen shown in Figure 154 on page 209.



X.25 Port - PLP Parameters 2/2

Port: 2304 Name: X252304A

PLP modulo: ☒ 8 ☐ 128

SVC Defaults

Packet window calling to called: 2 numerical [1-7]

Packet window called to calling: 2 numerical [1-7]

Packet size calling to called: 128 bytes

Packet size called to calling: 128 bytes

PLP Port Timers (seconds [1-255])

Restart timer (T20): 180 Reset timer (T22): 180

Call timer (T21): 200 Clear timer (T23): 180

PLP Port Maximum Retries (numerical [0-255])

Restart retry (R20): 1 Clear retry (R23): 1


Reset retry (R22): 1

OK Save as defaults Cancel Help

Figure 158. X.25 PLP Parameters 2/2

14. Select **IP parameters**.
15. Set MTU as desired.
16. Set IP Address=145.17.4.71. (This is the local IP address.)

17. Set IP Mask=255.255.255.0.
18. Select **OK**.



IP Over X.25 - Port Parameters

Port: 2304 Name: X252304I

General

Automatic reactivation? ☒ Yes ☐ No

Maximum transmission unit: 2058 bytes [576-4096]

☐ No idle timer Idle timer: 30 seconds [1-255]


Addresses

IP address:	Subnet mask:
145.17.4.71	255.255.255.0
145.17.4.71	255.255.255.0

Buttons: Add, Modify, Delete, OK, Save as defaults, Cancel, Help

Figure 159. X.25 Port Parameters

19. Select **APPN/IP Stations**.
20. Create an APPN station if desired.
21. Select **Add**.



X.25 Station Configuration

Port: 2304 Name: X252304I

Configure an X.25 Station

Network: ☒ APPN ☐ IP Station name: AP2304P ☒ PVC ☐ SVC

TOA/NPI? ☐ Yes ☒ No TOA: Network dependent {0} NPI: X.121 {3}

Remote DTE address: LCN: 2 PU type: ☐ 2 ☒ 2.1

Comments (optional):

X.25 Stations Already Configured


Network	Station name	Type	DTE address or LCN
IP	IP2304I3	SVC N/A	15552272
APPN	AP2304P	PVC 2.1	2

Buttons: Add, Modify, Copy..., Delete, Search..., Search next, DLC parameters..., APPN parameters..., IP parameters..., User facilities..., OK, Save as defaults, Cancel, Help

Figure 160. X.25 APPN Station Configuration

22. Create an IP station.
23. Set IP Station Name=IP2304S.
24. Set SVC.
25. Set TOA/NPI=YES.
26. Set Remote DTE address=155522721234567.

27. Select **Add**.
28. Select **OK**.



X.25 Station Configuration

Port: 2304 Name: X252304A

Configure an X.25 Station—

Network: ☒ APPN ☐ IP Station name: AP2304P ☒ PVC ☐ SVC **Add**

TOA/NPI? ☐ Yes ☒ No TOA: Network dependent (0) NPI: X.121 (3) **Modify**

Remote DTE address: LCN: 1 PU type: ☐ 2 ☒ 2.1 **Copy...**

Comments (optional):

X.25 Stations Already Configured—

Network	Station name	Type	DTE address or LCN
APPN	AP2304P	PVC 2.1	1
IP	IP2304S	SVC N/A	03 155522721234567

Delete **Search...**

Search next

DLC parameters...

APPN parameters...


IP parameters...

User facilities...

OK **Save as defaults** **Cancel** **Help**

Figure 161. X.25 IP Station Configuration

29. Select **IP parameters** for station IP2304S.
30. Set Remote IP Address=145.17.4.61.
31. Select **OK**.



IP Over X.25 - Station Parameters

Port: 2304 Name: X252304I

Station name: IP2304S

X.25 remote DTE address: 03 155522721234567

PVC logical channel number: N/A

General

Automatic reactivation? ☐ Yes ☒ No

Addresses

Remote IP address:

145.17.4.61 **Add**

145.17.4.61 **Modify**

145.17.4.61 **Delete**

OK **Cancel** **Help**

Figure 162. X.25 IP Station Parameters

32. Select **User Facilities**.
33. Set Add Calling DTE Address.

X.25 Station - SVC Call Requests User Facilities and Data

Port: 2304 Name: X252304I Network: IP
 Station name: IP2304S X.25 remote DTE address: 03 155522721234567

Protocol identifier: ☐ X'CB' ☐ X'CB' ☒ X'CC'

☐ Window size negotiation facility In: 2 Out: 2 num [1..7]
☐ Packet size negotiation facility In: 128 Out: 128 bytes
☐ Request reverse charging facility
☐ Request inter-network call redirection and deflection (ICRD) control facilities
 ICRD status selection: ☒ Prevention requested ☐ Allowance requested
☒ Add calling DTE address **Calling DTE address parameters...**
☐ Throughput class negotiation facility
 Format: ☒ Basic ☐ Extended In: 9.6 Out: 9.6 Kbps
☐ Closed user group selection facility ☐ With outgoing access selection
 Format: ☒ Basic ☐ Extended 0 numerical [0..99]

OK **Other UFD...** **Save as defaults** **Cancel** **Help**

Figure 163. X.25 User Facilities

34. Select **Calling DTE address parameters**.
35. Set Calling DTE address=144423041234567.
36. Select **OK**.

X.25 Station - SVC Calling DTE Address

Port: 2304 Name: X252304I Network: IP
 Station name: IP2304S
 X.25 remote DTE address: 03 155522721234567

Calling DTE Address—
 TOA: Network dependent (0) NPI: X.121 (3)
 Calling DTE address: 144423041234567

OK **Cancel** **Help**

Figure 164. X.25 Calling DTE Address Parameters

5.10 3746-9x0 IP Management

Figure 165 on page 215 shows the various components that play a role in the management of the 3746-9x0 IP functions. Depicted are:

Service processor (SP)

The service processor (SP) console gives an enhanced operator interface using the graphical user interface provided by IBM's OS/2. It provides maintenance and operator subsystem extended (MOSS-E) functions, loads the microcode of the 3746-9X0, provides a repository function for 3746-9X0 files such as the configuration data file-extended (CDF-E), and provides a repository function for the microcode running on the SP, NNP, and 3746-9x0.

The SP allows the operator to perform 3746 IP configuration, operation and management tasks with the controller configuration and management (CCM) program. For details see 5.10.1, “Controller Configuration and Management (CCM)” on page 216. The SP contains a TELNET client to access the 3746 IP TELNET configuration and operator functions (see 5.10.2, “TELNET Command Line Interface” on page 218) running on the NNP.




Figure 165. 3746-9x0 IP Management Components

Network node processor (NNP)

The network node processor (NNP) provides a TELNET server that enables TELNET clients to configure and operate the 3746 IP functions. For details see 5.10.2, “TELNET Command Line Interface” on page 218. The TELNET operator functions require TCP/IP connectivity between the TELNET client and the NNP. In addition, the NNP provides SNMP functions to enable SNMP management from a remote SNMP management station. For details see 5.11, “3746-9x0 SNMP Functions” on page 225.

3746-9x0

The 3746-9x0 IP router allows remote TELNET clients to access the NNP. Its SNMP relay function enables 3746 IP SNMP management from a remote SNMP management station. SNMP access from an

SNMP management station requires TCP/IP connectivity to the 3746 IP router. For details see 5.11, "3746-9x0 SNMP Functions" on page 225.

Operator station

All 3746-9x0 IP management functions are available from the SP. The SP functions can be accessed from the SP itself or from a remote station using DCAF. Remote access requires DCAF V1R3 or TME10 Remote Control V1R0 on the remote station and either TCP/IP, NetBIOS, or SNA connectivity to the SP.

Note: The dial-up connection to the SP can be used by DCAF/SNA only.

If only the 3746 IP operation and management functions are required on the remote station, one can TELNET into the NNP to access the 3746 IP command line TELNET interface (see 5.10, "3746-9x0 IP Management" on page 214). This requires a TELNET client function (VT100, VT200) and TCP/IP connectivity to the NNP.

Note: TELNET access from a remote SNA-attached station, other than DCAFin into the SP, requires AnyNet for OS/2 in addition to the TELNET client.

Network management station

The 3746-9x0 IP router provides an SNMP agent that can be accessed from a remote SNMP management station. SNMP management requires TCP/IP connectivity to the 3746-9x0 IP router.

Note: One can decide to combine the network management and remote TELNET operator function on a single workstation.

IBM Remote Service Facility (RSF)

The SP reports error codes to the IBM Remote Service Facility. Depending on the error reported IBM will access the SP to investigate and resolve the problem. To enable this service the SP is equipped with a dial-up modem. It is the customer's responsibility to provide a PSTN (telephone) attachment.

5.10.1 Controller Configuration and Management (CCM)

CCM provides a graphical user interface to configure and operate the 3746-9x0 IP functions. The configuration and operational functions are detailed in the following sections.

5.10.1.1 3746 IP CCM Configuration

CCM has been extended to provide 3746 IP configuration in addition to the 3746 NN/DLUR configuration functions introduced in 1995. An extensive description of how to use CCM for the 3746 IP functions, including many examples, can be found in Chapter 7, "Using CCM - IP Configuration Examples" on page 283.

CCM is available on the service processor and as a PC-based stand-alone version. The full set of CCM functions is available on the SP, while the stand-alone version is limited to preparing, saving, and exporting a 3746-9x0 configuration onto diskette. The 3746-9x0 diskette configuration can then be imported by CCM on the SP for further updating or activation.

Figure 166 on page 217 depicts the interaction of the stand-alone CCM, CCM on the SP, the NNP, and the 3746-9x0 microcode during the preparation and activation of a 3746-9x0 IP configuration file.




Figure 166. Preparing and Activating a 3746 IP Configuration

3746-9x0 configuration files can either be produced on the SP or on a stand-alone PC. A CCM Save operation on the SP results in a configuration file that contains a number of individual files packed together.

For details of the individual files see *Controller Configuration and Management User's Guide*, SH11-3081, Chapter 2 "About the Configuration Process", and Figure 191 on page 247.

During a CCM Activate operation, various files are copied to equivalent files on the NNP. As a result of the activation, the IP router code on all 3746 processors is then restarted.

Note: A CCM Save of a modified version of the active configuration results in an automatic Activate.

During a processor restart each 3746 processor retrieves its operational microcode from the SP and DLC configuration information from the configuration files on the NNP, and configures its ports for IP and/or APPN according to the DLC definitions.

During an IP restart each 3746 processor retrieves IP router configuration information from the files on the NNP and configures its IP functions according to the IP configuration definitions.

Note: Enabling a port for IP can only be done via CCM.

The IP IDF definitions that are used to configure the IP router can be viewed from CCM. To see the contents of this file select **Options -> View -> IP IDF**.

A restart of the IP routing functions is disruptive (for example, all routing information learned will be lost), but as IP is a connectionless protocol, external connection will not be interrupted.

Note: A CCM Activate will not impact the APPN NN functions unless the APPN configuration has been changed as well. For details of which action will impact other functions, see *Controller Configuration and Management User's Guide*, SH11-3081, Chapter 12 "Dynamic Activation and How it Affects your Network".

5.10.1.2 3746 IP Operation

The 3746 IP operator functions offered by CCM are discussed in 6.1, "Controller Configuration and Management (CCM)" on page 237.

5.10.2 TELNET Command Line Interface

The NNP TELNET server provides a command line interface to the 3746 IP operation and configuration functions. It provides a subset of services provided by CCM. Examples of the TELNET functions are described in 6.2, "Using TELNET - Operation and Configuration" on page 258. Only one TELNET connection can be active at a time.

Access to the 3746 IP TELNET operator and configuration services requires TELNETing into the NNP (see Figure 167). TELNET access to the NNP is provided from any station that has implemented the appropriate TELNET client function (for example, the TELNET client that is part of TCP/IP for OS/2). TCP/IP connectivity between the NNP and the TELNET client is required. To enable TELNET access from the SP, a TELNET client has been installed on the SP as well.




Figure 167. TELNET Operator and Configuration Interface

Note: The TELNET client functions on the SP are invoked by:

- Double-clicking on the 3746-9x0 icon
- Selecting the network node processor (NNP) management functions

- Selecting the IP commands management functions

After having connected to the TELNET server and having entered the appropriate user identifier and password (see discussion in 6.2.1, “OPCON Command Level” on page 258), 3746 IP TELNET operation and configuration can be started.

5.10.2.1 TELNET Command Levels

The line commands that are available via the TELNET operator and configurator facility are organized in a tree-like, hierarchical structure. Figure 168 depicts the various levels, the command prompts that can be expected within each level, and the commands to go from a specific level to the next lower level. The EXIT command lets you return to the next higher level. No shortcuts exist. For example, to go from GWCON to Config can only be done by returning to OPCON first.




Figure 168. TELNET Command Levels

An overview of the commands available within each command level can be obtained by entering ?. The three basic command levels are:

- OPCON

The OPCON level is entered after an operator connects to the TELNET operator and configurator facility. The OPCON command level is indicated by a command prompt ending with *.

For example:

```
RANGE 2048 2111 *
```

An overview of all OPCON commands available can be obtained by entering ?. For more details, see 5.10.2.2, “Operator Console (OPCON) Command Level” on page 221.

- GWCON

GWCON is the command level a TELNET operator enters after typing TALK 5 in OPCON. The GWCON command prompt ends either with + or >, as can be seen Figure 168 on page 219.

For example:

```
RANGE 2048 2111 +
```

GWCON offers a large number of commands to display the status of 3746 IP router components, display/reset router statistics, display/reset caches (for example, ARP), and diagnose networking problems. An overview of the commands available within each GWCON sublevel can be obtained by entering ?. To return to a parent level enter EXIT. For more details see 5.10.2.3, “Gateway Console (GWCON) Command Level” on page 222.

- Config

Config is the command level a TELNET operator enters after typing TALK 6 in OPCON. The Config command level is indicated by a command prompt ending with Config>.

For example:

```
Config>
```

The main function of Config is to enable configuration of the 3746 IP router functions. However, Config also offers commands to display the status of router components, list the active configuration, and diagnose networking problems. An overview of the commands available within each Config sublevel can be obtained by entering ?. To return to a parent level enter EXIT. For more details see 5.10.2.4, “Configuration (Config) Command Level” on page 224.

Note: Appendix A, “TELNET Line Command Overview” on page 329 gives an overview and a short description of all the 3746 IP TELNET line commands supported.

5.10.2.2 Operator Console (OPCON) Command Level

The TELNET OPCON command level is entered after TELNETing into the NNP. The main functions that can be done from OPCON are:

- Enter TALK 5 to enter GWCON
- Enter TALK 6 to enter Config
- Enter RESTART to restart the 3746 IP functions

An overview of all commands available can be obtained by entering ?.

It is important to understand that the actual OPCON command processor function does not reside on the NNP but is duplicated on each of the 3746-9x0 processors that perform IP routing functions. The NNP operates as a front end that interprets the OPCON commands, redirects them to the appropriate 3746-9x0 processor, and returns the response to the TELNET operator. The actual execution of the OPCON command is done on one of the 3746-9x0 processors.




Figure 169. OPCON Redirector Function

As mentioned earlier, the OPCON command processor process is duplicated on each 3746-9x0 processor. The OPCON prompt specifies the lowest and highest port address available on the 3746-9x0 processor to which the command will be redirected (and from which the response will be received).

For example:

```
RANGE 2048 - 2111* ---> OPCON commands redirected to CBSP
RANGE 2112 - 2207* ---> OPCON commands redirected to 2nd processor
RANGE 2208 - 2271* ---> OPCON commands redirected to 3rd processor
etc.
```

The RANGE command changes the 3746-9x0 processor to which successive OPCON commands will be redirected. For example, after entering RANGE 2144, the NNP will redirect successive OPCON commands to the second processor. The port addresses in the OPCON command prompt are updated to reflect the new 3746-9x0 processor on which the OPCON commands will be executed.

```
RANGE 2048 - 2111*range 2144
RANGE 2112 - 2207*
```

An exception to the rule that the OPCON commands are executed on a single 3746-9x0 processor is the RESTART command. This command will result in a restart of the 3746 IP routing function on all processors. The effect is similar to the activation of a new CCM configuration (see 5.10.1, “Controller Configuration and Management (CCM)” on page 216). During the restart of the IP routing function, the 3746-9x0 processors retrieve their configuration information from files on the NNP.

For a further discussion of OPCON, including some examples, see 6.2.1.3, “Basic OPCON Navigation” on page 261.

5.10.2.3 Gateway Console (GWCON) Command Level

Entering the TELNET GWCON command level enables you to issue a large number of commands to display the status of 3746 IP router components, obtain router statistics, and diagnose networking problems (for example, using PING or traceroute). GWCON commands are for display only and will not result in a changed configuration.

Within GWCON a number of command levels exist (see Figure 168 on page 219). The highest command level provides commands to display the more general router functions (for example, buffer sizes), while the lower level provides access to more specific functions, such as:

- Protocol IP, to obtain information on IP, RIP, Bootp, Access Controls, and IP filters
- Protocol OSPF, to obtain information on OSPF
- Protocol BGP, to obtain information on BGP
- Protocol ARP, to obtain information on ARP
- Feature BRS, to obtain information on BRS (bandwidth reservation)
- Network ‘xxxx’, to obtain information on port ‘xxxx’

The ? command lists all commands available within the present GWCON level. Entering a command and using ? as the first parameter results in an overview of the supported parameters for the command issued.

The RANGE command changes the 3746-9x0 processor to which successive GWCON commands will be redirected. For example, after entering RANGE 2144, the NNP will redirect successive GWCON commands to the second processor. The port addresses in the GWCON command prompt are updated to reflect the new 3746-9x0 processor on which the GWCON commands will be executed.

```
RANGE 2048 - 2111+range 2144
RANGE 2112 - 2207+
```

For a further discussion of GWCON including some examples, see 6.2.2, “GWCON Commands” on page 262.

5.10.2.4 Configuration (Config) Command Level

The main function of the commands available in the Config command level is 3746 IP router configuration. However, various commands exist to display the active configuration and perform problem determination.

Within Config a number of command levels exist (see Figure 168 on page 219). The highest command level provides commands to configure the more general functions, while the lower level provides access to more specific 3746 IP functions, such as:

- Protocol IP, to configure the IP, RIP, Bootp, access controls, and IP filters functions
- Protocol OSPF, to configure OSPF
- Protocol BGP, to configure BGP
- Protocol ARP, to configure ARP
- Feature BRS, to configure BRS (bandwidth reservation)
- Network ‘xxxx’, to configure hardware specifics on port ‘xxxx’

The ? command lists all commands available within the present Config level. Entering a command and using ? as the first parameter results in an overview of the supported parameters for the command issued.

The configuration commands are in general not immediately executed but instead, stored within configuration files kept on the NNP (see Figure 166 on page 217). The Config display commands have no effect on this file.

Note: As the Config commands are stored in the configuration files on the NNP when the user exits the Config command level. Be aware that your configuration commands may be lost if the TELNET connection is ended before returning to the OPCON command level.

Important

Terminate the TELNET connection in the OPCON command level only.

Storing the configuration commands in files on the NNP has a number of consequences:

- A 3746 IP router restart is required to activate the changes.
After the configuration commands have been entered, the operator needs to return to OPCON and enter RESTART to activate the changes. As discussed in

5.10.1, “Controller Configuration and Management (CCM)” on page 216, only the IP routing function will be restarted.

No RESTART is required for the following Config commands, as they are executed immediately:

- IP Config>add route (to add a static route)
- IP Config>delete route (to delete a static route)
- IP Config>add filter (to add a filter)
- IP Config>delete filter (to delete a filter)
- OSPF Config>add neighbor (to add an OSPF neighbor)
- OSPF Config>delete neighbor (to delete an OSPF neighbor)

For a description of the command syntax see Appendix A, “TELNET Line Command Overview” on page 329.

- CCM is required to enable ports for IP.

During restart the 3746-x90 processors retrieve the configuration files from the NNP and configure themselves accordingly. If IP has not been enabled for a port from CCM, then the configuration files on the NNP will not allow any configuration of IP at all on those ports.

TELNET configuration can be used to modify CCM port configuration. Definitions of ARP, Bootp, and the dynamic routing protocols do not require initial CCM configuration as they apply to the CBSP.

- TELNET and CCM configurations are not automatically synchronized.

Care should be taken when using the TELNET line command interface as the Config commands operate on the configuration files stored on the NNP. Changes to these files are not passed to the CCM files stored on the SP. Activating a configuration from CCM, will overwrite the files on the NNP. Any changes made to the files on the NNP via CONFIG since the last CCM activation are lost. It is the user's responsibility to synchronize the configuration files kept on SP and NNP.

For a further discussion of Config, including some examples, see Chapter 4, “IBM 3746-9x0 and IP Routing Protocols” on page 131.

5.11 3746-9x0 SNMP Functions

The 3746 IP router has implemented an SNMP agent to allow management from SNMP network management stations (for example, an AIX/6000 station running NetView/6000 for AIX). The 3746 IP SNMP agent supports the following MIBs:

- MIB II (described in RFC 1213)
- OSPF V2 MIB (described in RFC 1253)
- Token-ring MIB (described in RFC 1231)
- Frame relay MIB (described in RFC 1315)
- PPP MIB (described in RFC 1471)
- ESCON MIB (IBM proprietary)

The 3746 IP router provides SNMP V1 agent support. The SNMP GET, GET_NEXT, and TRAP commands are supported. The SNMP SET command is not supported. Every SNMP network management station equipped with the appropriate MIB support can retrieve this information.

To configure the 3746 IP SNMP configuration (community names, TRAP receivers, etc.) requires CCM configuration. SNMP configuration via the TELNET Config interface is not available. For a discussion of the CCM configuration see 7.12, “SNMP Parameter Configuration” on page 326.

5.11.1 3746-9x0 SNMP Relay

Figure 171 depicts the 3746 IP SNMP implementation. As can be seen two main components can be identified:

- SNMP relay code residing on the 3746-9x0
- SNMP agent code residing on the NNP

The SNMP relay function on the 3746-9x0 enables an SNMP network management station to send its SNMP requests to any of the IP ports of the 3746-9x0 IP router. By relaying SNMP traffic rather than forcing network stations to contact the NNP, the location of the SNMP agent is therefore transparent to the network management station.




Figure 171. 3746-9x0 IP SNMP Implementation

IP connectivity is required between the network management station and the 3746 IP router. The service LAN is used for the traffic between the 3746-9x0 and the NNP. Both 3746 IP and the NNP must be operational to enable SNMP access. Note that the SNMP traffic between 3746-9x0 and network management station is using UDP (port 161), while the SNMP relay function uses TCP transport.

5.11.2 Distributed Agents

The NNP SNMP agent uses the distributed programming interface (DPI V2) described in RFC 1592. To retrieve the information that SNMP network management stations are soliciting, the SNMP agent interfaces with a number of subagents. As most of the SNMP MIB information required is maintained on the 3746-9x0, the subagents themselves interact with processes running on the 3746-9x0 processors.

Figure 172 on page 228 details the subagents concept. Network management stations soliciting SNMP information access the 3746-9x0 which relays the requests to the NNP. The NNP accepts the SNMP request and forwards it to its appropriate subagent. The subagent decodes the ASN.1 format and interfaces with its counterpart in the 3746-9x0.

Note: For performance reasons the subagents retrieve whole tables that are cached to respond to successive SNMP requests.

The 3746 IP agent supports solicited and unsolicited SNMP traffic. Solicited data is sent in response to an SNMP query received from a network management station. Soliciting data requires read SNMP access on the 3746-9x0 IP routers. Figure 328 on page 327 depicts how the IP addresses and community names are set to enable read access. Unsolicited data is sent after errors have been detected. Unsolicited data is sent as TRAPs. Figure 329 on page 328 depicts how the IP addresses and community names are set for network management stations that should receive TRAPs.

5.11.3 SNMP Traps

Errors detected by the 3746-9x0 are categorized as external or internal errors. 3746-9x0 external errors are due to protocol violations, cabling problems, etc. 3746-9x0 internal errors are due to microcode and hardware problems.




Figure 172. Distributed SNMP Agents




Figure 173. 3746-9x0 IP SNMP TRAPs

Figure 173 depicts how the 3746-9x0 IP router generates TRAPs.

- External error

When an external error is detected the 3746-9x0 generates an alert to the NNP control point (CP). The NNP CP decides if the error is IP-related and invokes its RFCTRAP code. RFCTRAP generates the SNMP trap and hands it over to the SNMP agent. The latter forwards it to all SNMP network management stations that are configured to receive TRAPs.

- Internal error

When an internal error is detected the 3746-9x0 generates an alert to both NNP and SP. The processing of the alert sent to the NNP is equivalent to an external error and results in SNMP TRAPs being sent. For the alert received by the SP, configuration information is added before invoking RFCTRAP on the NNP and generating (a second) TRAP.

To handle the TRAPs on your network management station requires IBM's Router and Bridge Manager (RABM) that is presently available with Nways Campus Manager LAN which is shipped as a component of Nways Manager for AIX V1.2, and Nways Enterprise Manager. The Alert Manager within RABM is able to understand and convert the CMIP data that is contained within some of the TRAPs. For details on the information contained in the TRAPs see SA33-0175, *Alert Reference Guide*.

5.12 3746-9x0 Security

The following section explains how to configure IP filters and IP access control on the 3746.

5.12.1 3746-9x0 IP Filters

Figure 174 shows how to set an IP filter using CCM. To get to this screen do the following:

1. From the main CCM screen select **Configuration**.
2. Select **IP**.
3. Select **Filters**.




Figure 174. 3746-9x0 Configuring IP Filters

We set a filter to stop all traffic bound for subnet 9.0.0.0. This filters out all traffic where there is not a more specific route. Figure 175 on page 230 shows an example of using this filter. Traffic for destination IP address 9.10.2.2 is not filtered as there is a more specific route (9.10.0.0) than the filter.




Figure 175. 3746-9x0 IP Filters

5.12.2 3746-9x0 IP Access Controls

Figure 176 on page 231 shows how to set IP access control from CCM. To get to this screen do the following:

1. From the main CCM screen select **Configuration**.
2. Select **IP**.
3. Select **Access Controls**.

In the case shown, all traffic bound for IP address 9.100.0.0 with protocol number 17 will be discarded, as the entry is set to Deny/Exclusive. Note the non-deletable non-movable all-inclusive entry which is last in the list. This CCM screen also provides Move-up and Move-down buttons to allow the entries in the access control list to be ordered. This is necessary as the entries are processed from the top downwards in the order shown.

IP Access Controls

☒ Enable access control

Configure an Access Control Entry

Access control type: ☐ Permit/Inclusive (I) ☒ Deny/Exclusive (E)

Source: Network IP address 9.100.0.0 Mask address 255.255.0.0

Destination: 0.0.0.0 0.0.0.0

Protocol number: From 17 To 17 numerical [0-255]

Port number: 0 65535 numerical [0-65535]

Buttons: Add, Modify, Delete

Access Control Entries Already Configured

Type	Source IP/Mask addresses	Destination IP/Mask addresses
I	9.128.2.2 255.255.255.255	0.0.0.0 0.0.0.0
E	9.128.0.0 255.255.0.0	0.0.0.0 0.0.0.0
E	9.100.0.0 255.255.0.0	0.0.0.0 0.0.0.0
I	0.0.0.0 0.0.0.0	0.0.0.0 0.0.0.0

Buttons: OK, Move up, Move down, Save as defaults, Cancel, Help

Figure 176. 3746-9x0 Configuring IP Access Controls

The example shown in Figure 177 on page 232 shows IP traffic for three destination IP addresses arriving at the 3746. Traffic for IP address 9.128.2.2 is forwarded because of the first entry in the ACL (Inclusive - 9.128.2.2). If the first and second entries were reversed in order, then the entry Inclusive - 9.128.2.2 would prevent the second entry from ever matching a packet.

Traffic for IP address 9.128.2.3 is discarded because of the second entry in the ACL (Exclusive - 9.128.0.0). The first entry does not match this packet.

Traffic for IP address 9.100.1.1 is discarded because of the third entry in the ACL (Exclusive - 9.108.0.0 - Protocol=17). This IP packets is protocol 17 which is UDP. Other protocols to the same destination IP address would be forwarded as they would match the inclusive entry at the end of the list.




Figure 177. 3746-9x0 IP Access Controls

5.13 Changing Passwords

Information about passwords is distributed over two books in the original machine documentation. The *Service Processor Installation and Maintenance*, SY33-2109 describes the procedures used to change the passwords in detail. The *3745 Communication Controller Models A 3746 Nways Multiprotocol controller Models 900 and 950 - Planning Guide Part 2/3*, GA33-0457 provides more general information about passwords. The following is a brief description of how to change the passwords. Table 6 on page 233 displays the password modes.

Table 6. MOSS-E Password Modes					
	3746 Nways Multiprotocol Controller Menus		Service Processor Menu		
	Operator Functions	Maint. Functions	Operator Functions	Maint. Functions	Password Mgt.
Controller customer	X				
Controller maintenance	X	X			
Service processor customer	X		X		
Service processor maintenance	X	X	X	X	
Management password					X
Telnet password					X
Legend: Mgt Management Maint Maintenance					

To change the passwords select **Operations Management** on the Service Processor Menu (see Table 6).




Figure 178. Service Processor Menu

After that you must enter the Management Password (see Figure 179 on page 234). The default for this password is *ibm3745*.




Figure 179. Password Management

In the next screen you can decide which password you want to change. Select the required password and click **OK**.




Figure 180. Password Management Selection Menu

Figure 181, Figure 182 on page 235 and Figure 183 on page 235 show the displayed screen to change the specific password(s).





Figure 181. MOSS-E View Password



Management Password


Password: IBM3745

Failed attempts: 0

Attempt threshold: 99

Buttons: Apply, Reset, Reset failed attempt counter, Cancel, Help

Figure 182. Management Password



CCM/Telnet User Profiles Management

☐ Enable CCM/telnet user profiles

Configure a User Profile

Userid: Password:

Access type: ☐ No access ☐ View only ☒ All

Access on which 3746-9X0?

☐ 3746-9X0 S/N: 57-91805 Name: ITS0

☐ Not used

User Profiles Already Configured

Userid	Password	Access type	3746-9X0 S/N

Buttons: Add, Modify, Delete, OK, Cancel, Help

Figure 183. CCM/Telnet User Profiles Management

To return from the password management function click **Cancel** in the Password Management Selection Menu (see Figure 180 on page 234).

5.14 3746 IP Tuning Recommendations

The following is a short list of IP tuning recommendations. The information here is not meant to be complete but it is information which may very quickly have a positive effect on IP performance.

For more information on TCP/IP tuning, refer to the *IBM TCP/IP Performance and Tuning Guide*, SC31-7188.

- Set TCPIP despatching priority approximately the same as VTAM, and the other TCP/IP address spaces a little below TCPIP.

- Specify a large DATABUFFERPOOLSIZE for better performance (set it to 32768). (NETSTAT POOLSIZE provides useful information for how buffers and control blocks are being used).
- Make client and server TCP window sizes equal and as large as possible
- Make sure there's nothing in the network path which reduces the TCP packet size. It's a case of lowest common denominator and if the MTU packet size is less than 1500, then performance is going to suffer.
- Use TCP/IP for MVS V3R2 or above, or the OS/390 TCP/IP stack.
- Use large file block sizes, for example, 0.5 MB of a DASD track.
- Use DELAY=0 for each ESCON IP station.

The delay is set from CCM. For each ESCON IP station select **DLC Parameters**. The resulting dialog allows the user to set delay, channel adapter slowdown (CASDL), and the attention timer (TIMEOUT) values.

- Increase number of buffers if necessary by monitoring POOLSIZE as follows:
 - Start the TCP/IP system.
 - Issue a NETSTAT POOLSIZE and note the maximum buffers used before you start your application.
 - Run your application.
 - Issue a new NETSTAT POOLSIZE and note the maximum buffers now.
 - The difference is the number of buffers your application used.

- Set the DEVICE statement to 15 15 4096 4096.

The product of read_buffers (15) * read_size (4096) must not exceed 65535.

- In the TCP/IP profile set the following parameters:
 - DATABUFFERPOOLSIZE 160 32760
 - LARGEENVEELOPOOLSIZE 50 32760
- In the GATEWAY statement check that the packet size is set at 4092.

This should be set to the value of read_size - 4 (4096 - 4).

- If possible, define a direct route as follows:
 - xx.xx.x.x = CDL1 4092 HOST
- Define a default route as follows:
 - DEFAULTNET xx.xx.xx.x CDL1 4092 0

Chapter 6. 3746-9x0 IP Operation and Configuration

Two methods are available to the user to operate and configure the 3746-9x0 IP router functions. The first is using the Controller Configuration and Management (CCM) interface on the service processor. The second method is using the 3746 IP TELNET line commands service that can be accessed via the Network Node Processor (NNP).

For details on IP-specific CCM management, see 6.1, "Controller Configuration and Management (CCM)." CCM configuration functions are detailed in Chapter 7, "Using CCM - IP Configuration Examples" on page 283. For a description of the TELNET operator and configurator functions, see 6.2, "Using TELNET - Operation and Configuration" on page 258.

Note: It is advised that you read 5.10.1, "Controller Configuration and Management (CCM)" on page 216 and 5.10.2, "TELNET Command Line Interface" on page 218 first, before reading the following sections.

6.1 Controller Configuration and Management (CCM)

CCM on the 3746 service processor provides the ability to configure and manage the 3746 APPN and IP resources using a graphical user interface (GUI), as opposed to the line command interface provided by the NNP TELNET server.

6.1.1 Welcome to the CCM

The IBM Controller Configuration and Management application (CCM) is designed to help you configure and manage an IBM 3746 communication controller and its associated network resources.

When you configure your controller and its resources, the CCM creates a configuration file, referred to as the **3746 controller configuration file**. Using the CCM, you can create several configurations.

The CCM runs under the control of the OS/2* and features a presentation manager graphical user interface, where you can perform a wide range of tasks.

The CCM tasks are divided into two main categories:

Configuration

This is for defining configuration parameters such as coupler type, mode, class of service, transmission group, and others. When a group of configuration parameters has been defined, it can be saved to file on disk. This file can then be immediately activated for use by the network, or it can be saved for later use. Configuration parameters are defined by specifying values in CCM windows.

Management

This involves viewing operational information about the currently defined network resources, and activating or deactivating network resources to maintain optimal network performance. Tasks requesting network resource information use commands that only specify the resource address.

For more information on using the CCM, refer to:

- *3746 Communication Controller Models A, 3746 Nways Multiprotocol Controller Models 900 and 950: Planning Guide*, GA33-0457
- CCM online help
- *3746 Nways Multiprotocol Controller, Models 900 and 950: Controller Configuration and Management User's Guide*, SH11-3081

6.1.2 Operating Environments

The CCM can be used on either:

- The service processor, where it is accessed via the MOSS-E. (The service processor may be accessed via a Distributed Console Access Facility (DCAF) remote console.) This environment is referred to as the ***service processor CCM version***.
- A stand-alone workstation. This environment is referred to as the ***stand-alone CCM version***.

6.1.2.1 Service Processor Environment

With the CCM installed and running on the service processor, you access it through the MOSS-E user interface. In this type of installation, both the configuration and the management functions can be used.

6.1.2.2 Stand-Alone Environment

When the CCM is running in the stand-alone environment, the management part of the application and the coupler with the 2080 address are not available for use (they are 'greyed-out').

However the configuration part of the application is fully available for configuring the controller and its resources before your machines arrive.

If several controllers are operating on the network, a good strategy is to configure all controllers from a centralized location, using the stand-alone CCM. The configurations can then be sent (exported) to each service processor when complete.

Minimum Hardware and Software Requirements: The minimum requirements for workstations running the stand-alone version of the CCM are:

- 80486 microprocessor or higher.
- 40 MB of hard disk space free.
- VGA display (for example, an IBM 8515 Color Display or equivalent).
- 24 MB of virtual memory. The actual amount of virtual memory needed depends on the size of the configuration (does it have tens of lines or hundreds of lines).
- Mouse.
- 3.5-inch diskette drive.
- IBM Operating System/2 (OS/2), version 2.1 or higher.

6.1.3 Installing CCM

This section explains the CCM driver levels and describes the MOSS-E and stand-alone installation procedures.

6.1.3.1 CCM Levels

Table 7 shows a list of ECA levels with the corresponding CCM EC levels.

<i>Table 7 (Page 1 of 2). Functions Supported by Each CCM Level</i>			
ECA Number	Microcode EC Level (See note 1)	CCM Version (APPN BLP EC Level)	Functions Supported
144 146	D22560A D22560D	D22561 D22561	APPN configuration and management
155	D46100	D22571	The above functions plus: • IP • Frame relay over APPN • HPR/ANR • IP configuration for ESCON and token ring.
157 159 159	D46120 D46120A D46120B	D46121 D46121.005 D46121.010	The above functions plus: • IP management • PPP • Frame relay over IP • HPR/RTP • Second expansion enclosure • A CCM password • NetView Performance Monitor (NPM).
167	D46130	D46131.000	The above functions plus: • HRP/RTP and ARB on token ring, Ethernet, SDLC, frame relay, and ESCON • APPN/ISR (5000 PUs + 15 000 sessions) • Display of 3746 EC/ECA microcode level • CDF-E checking
170 (See note 2)	D46130I	D46131.024	The above functions plus: • X.25 support on 3746 (under NNP control) • HPR MLTG on token ring, SDLC, frame relay, and ESCON • BRS for 3746 PPP lines • FRFH on 3746 lines • CIR on 3746 frame-relay lines.

Table 7 (Page 2 of 2). Functions Supported by Each CCM Level

ECA Number	Microcode EC Level (See note 1)	CCM Version (APPN BLPU EC Level)	Functions Supported
175 (See note 2)	F12380	F12381.000	The above functions plus: • APING function added (like IP PING, but for APPN traffic) • ESCON adapter re-IML is now optional when activating a configuration • Increased management of remote IP addresses • Make DLCI copies function added • IP access controls now available at port level.
Notes: 1. This is the minimum functional EC level as of the General Availability date of the CCM level. 2. This ECA must not be ordered, the corresponding microcode is automatically shipped with any features requiring this level. In particular, the 3746 Extended Functions 1 (feature number 5800) must be ordered to operate any the functions listed. 3. Legend: BLPU build logical program unit EC engineering change ECA engineering change announcement			

Additional, level-specific information is available in the README file that comes with the stand-alone CCM.

Level compatibility

Different levels of CCM are upwardly compatible only. For this reason, a configuration that has been generated using CCM at one level can only be exported to a CCM at the same or higher level.

IBM recommends that the same level of the CCM should be maintained in the service processor and the stand-alone environment. This is important if configurations generated on the stand-alone CCM are to be exported to a CCM running in a service processor.

Details about the CCM level are displayed in the **Product Information** window.

6.1.3.2 Viewing the CCM Product Information

You can view information about this release of the CCM: its version, EC level, and the date of general availability of this EC version. To view the product information:

Main window **Help** menu → **Product information**




Figure 184. Main Window Help Menu

6.1.3.3 Installing CCM in Service Processor Environment

Installation of the CCM in the MOSS-E is a task which is performed by an IBM customer engineer. The CCM is installed when the MOSS-E is installed.

6.1.3.4 Installing CCM in Stand-Alone Environment

The IBM customer engineer creates the CCM installation diskettes that you use for installing the CCM.

Before installing CCM

Ensure that your workstation has the correct hardware and software requirements (see page 238).

The installation procedure is in the README files that comes with the CCM and may be different according to the level of the CCM that is being installed.

6.1.3.5 CCM Password Protection from MOSS-E

The CCM on a service processor can be protected by a password using the MOSS-E CCM/Telnet **User Profiles Management** function in the **Manage Passwords** menu (see 5.13, "Changing Passwords" on page 232).

6.1.3.6 Starting the Service Processor CCM

This section walks through the Controller Configuration and Management (CCM) panels on the service processor to define the connections. Connections have been verified as far as permitted by the hardware and software available at the time of writing this publication.




Figure 185. MOSS-E View

Figure 185 depicts the MOSS-E view of the 3746-900. Double-click on the 3746-9x0 icon to access its menu functions.




Figure 186. MOSS-E 3746-9x0 Menu

Figure 186 illustrates the 3746-9x0 Menu options available. CCM functions are accessed by selecting the **Network Node Processor (NNP) Management** menu option.




Figure 187. Selecting CCM




Figure 188. CCM Logo

Figure 187 on page 242 illustrates the NNP Management functions. Double-click on **Controller Configuration and Management (CCM)** to start CCM. Figure 188 illustrates the CCM logo displayed upon invocation.




Figure 189. CCM Main Configuration Menu

Figure 189 results after starting CCM. The active configuration is shown.

6.1.3.7 Starting the Stand-Alone CCM

Note: When the CCM is running in the stand-alone environment, the management part of the application and the 2048 coupler are not available for use (they are 'greyed-out').

To start the stand-alone CCM, you can use either the mouse or the keyboard.

Using the Mouse



Step 1. Double-click on the CCM folder icon CCM



Step 2. Double-click on the CCM icon CCM to start the application.

Using the Keyboard

Step 1. Open an OS/2 window.

Step 2. Type CCM and press **Enter**.

Stopping and Exiting from CCM: To stop the CCM:

Main window **File** → **Exit**

6.1.4 Becoming Familiar with the CCM User Interface

The CCM provides a graphical user interface with which you will quickly become familiar.

When you start the CCM, the main window is displayed (see Figure 190).




Figure 190. CCM Main Window

As shown in Figure 190, the main window includes the following features:

- Menu bar
 - Five pull-down menus are available from the menu bar:
 - File
 - Configuration
 - Management
 - Options
 - Help.
- Header information

- Coupler icons
- Information area.

6.1.4.1 Header Information

The header information is dynamically updated, and shows the following:

- The active configuration (if one exists), not available in stand-alone CCM
- The opened configuration (if one has been opened)
- The opened reference CDF-E file (if one exists in the opened configuration).

6.1.4.2 Coupler Icons

The coupler icons represent a schematic view of the couplers in the basic enclosure and expansion enclosures within the controller.

The lower row of coupler icons represents the 10 coupler slot addresses in the basic enclosure, which can hold up to five adapters.

The middle row of coupler icons represents the 12 coupler slot addresses in the first expansion enclosure, which can hold up to six adapters.

The upper row of coupler icons represents the 12 coupler slot addresses in the second expansion enclosure, which can hold up to six adapters.

An adapter consists of one processor connected to one or two couplers: each adapter has a pair of coupler icons. Adapters configured for ESCON channels only use one of a coupler pair.

The shape and color of the coupler icons give information about the couplers they represent.

Icon Shape There are two shapes for the coupler icons:



Means that the coupler has been configured in the CCM.



Means that the coupler has not been configured in the CCM.

For example, if the first coupler configured on a processor is a TIC3, the second slot on the processor is automatically labeled as an unconfigured TIC3, since only a TIC3 can be installed in the second slot.

In the same manner, if a LIC (LIC11 or LIC12) is configured on a processor, the other slot is labeled as a LIC.

Blue and White Icons These colors indicate the status of a coupler after a comparison has been made between the actual, physical configuration of the coupler slot (as given in the reference CDF-E file) and the CCM configuration of the coupler slot:

Blue Means that, when the CDF-E file was compared to the CCM configuration file, no discrepancies were found between the two for that coupler slot.

Note: While a coupler is blue, its DLC type cannot be changed in the CCM configuration unless the configuration is cleared.

White Means one of the following:

- The reference CDF-E/CCM configuration comparison has not been made
- During the comparison, the CCM coupler DLC type **was not** found in the reference CDF-E file.
- During the comparison, the CCM coupler DLC type **was** found in the reference CDF-E file, but there are configuration discrepancies that would prevent the coupler from operating properly.

Grey Icons



Means that the coupler can be neither selected nor configured because it is the:

- Second, unused slot in an ESCON adapter.
- 2048 slot, which is never available because of the amount of traffic for the NNP and service processor handled by the 2080 TIC3 attached to the service LAN. The 2048 and 2080 slots are for the CBSP (type 2 or 3).
- 2112 slot, this will be grey if the 3746-900 is attached to a dual-CCU 3745.

Note: In the 3746-900, the 2048 is used to connect the 3745 to the 3746. This connection does not handle (route) a large amount of traffic, most of it just passes through and is handled by other 3746 processors.

6.1.4.3 Information Area

The information area is located at the bottom of the main window and displays navigation and processing status information.

6.1.4.4 Working in the Main Window

This section briefly explains how to work with the coupler icons and the menus in the main CCM window.

6.1.4.5 Working with the Main Window Menu Choices

In a stand-alone CCM version with no configuration currently opened, you can:

- Create a new configuration: **File** menu → **New**
- Open an existing configuration: **File** menu → **Open**
- Import a configuration: **File** menu → **Import**

Note: All choices on the **Options** and **Help** menus are also available.

6.1.4.6 Working with the Coupler Icons

Select a coupler icon to begin configuring a coupler. If you are using a mouse, double-click on the icon for the coupler you want to configure. If you prefer to work with the keyboard, use the arrow keys to highlight the desired coupler and press **Enter**.

If the coupler type has already been defined, you go directly to the window needed to continue the configuration process.

If the coupler type has not been defined, CCM opens the **Coupler type** window, where you must specify the coupler type. CCM then opens the window needed to continue the configuration process.

If the configuration is new, the first time you select a the coupler, CCM opens the **3745/3746 Parameters** window, where you must specify the 3746 model used and give information about the 3745 if a 3746 Model 900 is being used. CCM then opens the **Coupler type** window to continue the configuration process.

6.1.4.7 Working from the Pull-Down Menus

For some pull-down menu choices, you do not have to select a coupler icon, instead just select the option from the menu. For example, any of the **File** pull-down menu choices can be selected without first selecting a coupler icon.

6.1.5 About the Configuration Process

The CCM enables you to configure the controller and all the associated resources it uses for handling network traffic.

6.1.5.1 An Easier Way to Configure

The CCM is designed to provide a much simpler method of configuring the controller and its resources, when compared with the NCP generation process.

About 80% of the parameters have predefined default values. These values can be used, or modified and saved as new default values if required. This saves time and effort when several identical lines, ports, or stations are being configured.




Figure 191. Files Created by the CCM during the Configuration Process

Configuration files can also be copied or exported (or printed as listings). The CCM ensures that the configuration is internally consistent by dynamically cross-checking the validity of parameter values while you are entering them.

This dynamic checking allows CCM to help you, when you choose a value for a parameter, by immediately disabling (greying-out) all the other CCM parameters that are:

- No longer relevant to the configuration you are working on.

For example, if you start to configure an ESCON port for a SNA/subarea network, the APPN and IP name parameters are greyed-out as they do not apply to a SNA/subarea ESCON Channel.

- Automatically selected by the CCM and cannot be changed.

For example, for a serial line port, if you choose the DLC as SDLC, the network parameter is automatically set to APPN and the set of choices is greyed-out. But, if you choose the PPP DLC, the network is automatically set to IP and, again, the set of choices is greyed-out.

During the configuration process, the CCM creates a set of output files which are then compressed into a single file known as the 3746 configuration file (the .CCM binary file, see Figure 191 on page 247). Table 8 contains details about the CCM files produced.

Table 8. Files Created by the CCM during the Configuration Process		
File Extension	Details	
.CMD	Name	SNMP Definition File
	Contents	SNMP definitions
	Destination	Network Node Processor (NNP)
.NDF	Name:	Network Definition File
	Contents:	APPN and IP resource configurations
	Destination:	Network Node Processor (NNP)
.IDF	Name:	Internet Definition File
	Contents:	IP resource data
	Destination:	Processors that handle IP traffic
.IOC	Name:	I/O Configuration Program file
	Contents:	Defines the ESCON channel paths
	Destination:	Destination: Host (IOCP or the MVS Hardware Configuration Definition (HCD) tool) Note: The CCM produces this file as output from a 3746 configuration file to be used as input for the host.
.NCP	Name:	Network Control Program (NCP) file
	Contents:	ESCON definitions for NCP
	Destination:	Host (NCPGEN) Note: The CCM produces this file as output from a 3746 configuration file to be used as input for the host.
.ESC	Name:	ESCON Definition File
	Contents:	SNA/subarea, APPN, and IP ESCON definitions
	Destination:	Used to configure ESCON processors
.CCM	Name:	CCM Configuration File
	Contents:	Complete CCM configuration (compressed) with all the above files and others
	Destination:	Hard disk that contains the CCM program

6.1.5.2 What You Can Do with a CCM Configuration

With the CCM, configuration files can be created, modified, copied, imported, exported, and activated as required.

Creating Configuration Files: Different configuration files can be created for different controller configurations and environments (but only a single configuration file can be active at a given time).

You can create a configuration file in the service processor CCM version or in stand-alone environment.

Before starting the initial configuration ensure you have available:

- The hardware configuration worksheets, which are located in the *3746 Communication Controller Models A, 3746 Nways Multiprotocol Controller Models 900 and 950: Planning Guide*, GA33-0457, can be used to keep a record of the controller hardware topology, including details of coupler position and type.

Modifying Configuration Files: If resources are changed (for example, if a coupler is added or a coupler is replaced with one of a different type) the configuration file must also be updated.

This can be done in the service processor CCM version or in a stand-alone environment.

After modifying a configuration, you can activate the changes in the configuration either:

- Immediately, using the CCM Dynamic Configuration Update function on individual ports and stations without disrupting the rest of the network.
- Later, by activating the whole configuration.

Before modifying a configuration:

- Know the file name of the configuration to be modified.
- Have the hardware configuration worksheet, if the hardware topology of the machine has changed.
- Have the parameter worksheets with the details of the changes to be made in the configuration.

Copy a Configuration: You can make one or more copies of a given configuration by saving its file under different names.

This is used, for example, if several controllers on the network have similar configurations and a "master" configuration contains most of definitions needed by all these controllers. This master configuration can be changed as needed for an individual controller and saved under a unique name that corresponds to this controller. This can be repeated for each of the other controllers, giving you a group of configurations each customized for a specific controller.

This method can also be used to produce several configurations for the same controller. For example, to handle traffic over a coupler differently at night, the active configuration could change at 20:00 to the night version and change again at 07:00 to the day version.

Export/Import a Configuration: If you configure in the stand-alone environment, or in the service processor for a controller not attached to the service processor, the configuration file must be exported from the CCM to a diskette and then imported onto the service processor hard disk of the destination controller using its CCM and MOSS-E. You can also import a CCM configuration file into directories other than the default directory.

Activate a Configuration: This must be done in the service processor CCM version CCM.

Only one configuration can be active at a time.

6.1.5.3 Configuring the Controller

To configure a controller, you define the parameters for the:

Controller itself:

- Controller frame information
- Its network focal point
- As a dependent LU requester (DLUR)
- Its mode of CCU operation (for a 3746-900)
- Class of service (COS) for its traffic
- Communications protocol

Controller resources:

- The couplers

To configure a coupler, you define parameters for its ports and stations.

The controller and its resources must be configured when they are first installed and when modifications are made to the network.

6.1.5.4 Configuration Creation in Different Environments

The procedure for creating a configuration depends on the environment in which you are working (service processor or stand-alone).

Note: In the following figures, the activation step has been included to show the difference between the two environments.

In the Service Processor CCM Version: If you are using the service processor CCM version, follow the steps shown in Figure 192 on page 251.




Figure 192. Creating a Configuration in the Service Processor CCM Version

In the Stand-Alone CCM Version: If you are using the stand-alone CCM version, follow the steps shown in Figure 193.




Figure 193. Creating a Configuration in a Stand-Alone CCM Version

6.1.5.5 Creating a New CCM File




Figure 194. File Options

Figure 194 illustrates how a new configuration can be started. Select **File** from the pull-down menu option and select the **New** option.




Figure 195. Configuration Description

Figure 195 results. Enter the description of the configuration in the Name field. Add some useful comments for documentation purposes. If a previous ESCON configuration has been generated with the ESCON Generation Assistant (EGA) program, then the information can be imported in your new configuration by selecting **Import ESCON SBS file created with EGA**.




Figure 196. Opened Configuration

Figure 196 depicts the resulting CCM panel display when a new file or existing file is opened. The opened configuration file and the active configuration file names are both indicated. In our illustration, a new file called SG24-4845-01 has been opened.

6.1.5.6 Saving a CCM File

Once the configuration is completed the CCM file can be saved. Before the file can be saved, CCM requires that the machine Network Node identifier (see Figure 198 on page 254) and machine type (see Figure 200 on page 255) are specified.




Figure 197. NN FP/DLUR... Configuration Option

The Network Node identifier and other applicable parameters become available from the NN/FP/DLUR option in the Configuration pull-down menu (see Figure 197).

Network Node/Focal Point/Dependent LU Requester Parameters

Network Node and Focal Point Parameters

Network identifier Control point name
 Network Node (NN): USIMBRA . NN041A
 Comments (optional):
 Network identifier Control point name
 Network management Focal Point: .
 HPR support: HPR base (ANR)

NN characteristics... Backup focal point... RTP parameters...

Dependent LU Requester (DLUR) Parameters

Network identifier Control point name
 Primary dependent LU server (DLUS): .
 Backup DLUS? ☐ Yes ☒ No .
 Waiting time before short retry: 10 seconds [- 1] - 120
 Waiting time before long retry: 30 seconds [30 - 1200]

OK Cancel Help

Figure 198. Network Node/Focal Point/Dependent LU Requester Parameters

Figure 198 illustrates the Network Node details that are applicable. The Network Node (NN) identifier and Control point name are mandatory parameters (even if the 3746-9x0 is used for IP functions only). The Focal Point, its backup, HPR support along with DLUR and DLUS parameters are optional and should only be entered if the 3746 NN/DLUR function is used as well.




Figure 199. 3745/3746... Configuration Option

The 3746-9x0 machine type and mode of operation can be entered after selecting the 3745/3746 option in the Configuration pull-down menu (see Figure 199). Note that this option can only be selected after at least one 3746-9x0 coupler has been configured.




Figure 200. 3745/3746 Parameters

Figure 201 on page 256 illustrates the 3745/3746 parameters. Enter the correct 3746 model and, if the model is a 3746-900, indicate the 3745 mode of operation. Note that this panel will be displayed upon the first coupler defined and, thereafter, can be accessed from the Configuration pull-down menu. The panel is displayed during configuration of the first 3746-9x0 coupler to allow CCM to identify the CBC positions and allow proper ESCON configuration.




Figure 201. Saving a Configuration File

Figure 201 illustrates how a CCM configuration file can be saved using the **Save** option from the File pull-down menu.

6.1.5.7 IP Management from CCM

The IP management functions allow you to view information about ports, obtain details about the active configuration, and display the routing table. Traceroute and PING functions are available to verify connectivity. An overview of the CCM IP-specific functions is depicted in Figure 202 on page 257. OSPF and BGP functions are detailed in Figure 203 on page 257 and Figure 204 on page 258, respectively.

Note: IP management commands from CCM also use the OPCON process in the NNP. We found sometimes that if a user is logged on to the NNP via TELNET, then CCM cannot issue IP management commands from the SP.




Figure 202. CCM Management IP Specifics Commands Display




Figure 203. CCM Management IP OSPF Specifics Commands Display




Figure 204. CCM Management IP BGP Specifics Commands Display

6.2 Using TELNET - Operation and Configuration

An alternative to the CCM configuration and management services is to use the 3746-9x0 IP TELNET operator and configuration interface. As discussed in 5.10.2, “TELNET Command Line Interface” on page 218, a TELNET server has been implemented on the NNP for this purpose.

The TELNET interface provides three command levels: OPCODE, GWCON, and CONFIG. A description of each and a number of examples are given in the following sections.

Note: Appendix A, “TELNET Line Command Overview” on page 329 gives an overview and a short description of all the 3746 IP TELNET line commands supported.

6.2.1 OPCODE Command Level

When connecting to the NNP TELNET server, the user automatically enters the OPCODE command level.

6.2.1.1 What Is OPCODE

The Operator Console Process (OPCODE) is the root level process of the 3746 IP TELNET operation and configuration service. Using OPCODE commands, you can:

- List all available OPCODE commands
- Display the status of the IP router processes
- Enter the Gateway Console Process (GWCON)
- Enter the Configuration Process (CONFIG)
- Display IP router memory usage

- Restart the 3746-9x0 IP router

6.2.1.2 Connecting to the NNP TELNET Server

The NNP TELNET server can be accessed from each workstation that has TCP/IP connectivity to the NNP and is equipped with the appropriate TELNET client (VT100, VT220). Of course you can establish a telnet session from the SP. Select **Network Node Processor (NNP) Management** in the 3746-9x0 Menu. The corresponding subtree appears. In this subtree select **IP Commands** (see Figure 205).




Figure 205. Selecting IP Commands

In Figure 206 you can see the result. From then on there is no difference between using SP Telnet or any other Telnet client.




Figure 206. Telnet Session

After connecting to the TELNET server, one automatically enters the OPCON command level.

The following shows an example using the TELNET client on a PC workstation using TCP/IP for OS/2:

1. Telnet from the workstation to the NNP. Indicate hereby the NNP IP or if your installation has implemented a domain name server, the domain name assigned to the NNP.
2. Once connected, enter the login identifier at the login: prompt. The default identifier is NNPIP.
3. After the correct login identifier has been entered, enter the password at the password: prompt. The default password is 37469x0a.

Figure 207 illustrates the TELNET login sequence. Note that OPCON uses * as a command prompt suffix.

```
C:\>TELNET 192.168.152.2  
  
OS/2 Version 2.1 (CA191805)  
  
OS/2 2.11 (CA191805)  
login: nnpip  
password: 37469x0a  
login successful.  
Checking Status  
CA191805 is active.  
Launching IP Commands  
RANGE 2048-2111 *
```

Figure 207. NNP TELNET Logon Sequence

IP management commands from CCM also use the OPCON process in the NNP. If a user is logged on to the NNP via TELNET, then CCM cannot issue IP management commands from the SP.

Access to the TELNET server may fail for a number of reasons. Two reasons are worth mentioning.

Only one TELNET connection can be active at a time. The problem illustrated in Figure 208 on page 261 occurs when attempting to start a second TELNET connection.


```

C:\>TELNET 192.168.152.2

OS/2 Version 2.1 (CA191805)

OS/2 2.11 (CA191805)
login: nnpip
password: 37469x0a
login successful.
Checking Status
CA191805 is active.
Launching IP Commands
iptotnu queue create failed rc=332
An other telnet user is probably logged already...
Disconnecting ...

```

Figure 208. NNP TELNET Second Session Logon Sequence Failure

Another reason for a TELNET access failure is when no TELNET connection can be established due to connectivity problems. It is possible that the client station cannot reach the NNP due to network failures, access control or filters defined on the 3746 IP router, or an incorrect IP address. Figure 209 illustrates the result when the target IP address exists but no TELNET server is supported on this port. In our case we addressed the 3746-9x0 IP router by its service LAN attachment instead of the NNP IP address.

```

C:\>TELNET 192.168.152.3

===== Telnet Message =====
=                                                                           =
= Telnet connect: Connection refused                                     =
= Hit any key to exit _                                                =
=                                                                           =
=                                                                           =
=====

```

Figure 209. NNP TELNET Invalid Address Logon Sequence Failure

To minimize the chance of failure, verify the IP connectivity with an IP PING command.

6.2.1.3 Basic OPCON Navigation

The following provides you with the basic navigation through the levels of OPCON:

- ? will list the available commands at the current level.
- TALK 5 connects you to the GWCON command level allowing for monitoring functions.
- TALK 6 connects you to the CONFIG command level allowing for configuration functions.
- EXIT returns you to the previous command level.
- Logout or [CTRL] C will end the TELNET connection.

Note: CONFIG commands are not be stored on the NNP until you retur to the OPCON command level. Therefore, if the TELNET connection is ended before returning, your configuration commands may be lost.

Important

Therefore, terminate the TELNET connection in the OPCON command level only.

Figure 210 illustrates the basic navigation commands at the OPCON command level.

```
RANGE 2048-2111 *?  
RANGE 3746-port-number  
MEMORY statistics  
RESTART  
STATUS of process(es)  
TALK to process
```

Figure 210. Basic OPCON Commands

Exit OPCON (end the TELNET connection) with [CTRL] C:

```
RANGE 2048-2111 *[CTRL] C  
RANGE 2048-2111 *  
Disconnecting ...  
C:\>
```

Figure 211. Exit OPCON

6.2.2 GWCON Commands

The Gateway Console Process (GWCON) commands enable you to display the status of IP resources, display routing statistics, and diagnose network problems. GWCON is accessed with the OPCON TALK 5 command, as illustrated in Figure 212. Note that the command prompt suffix changes to +.

```
RANGE 2048-2111 *talk 5  
RANGE 2048-2111 +
```

Figure 212. GWCON Command Level

The ? command lists the available commands (see Figure 213 on page 263). Some of the GWCON commands we used during our testing are explained in 6.3, “TELNET Line Command Examples” on page 266.

```

RANGE 2048-2111 +?

RANGE 3746-port-number
BUFFER statistics
CLEAR statistics
CONFIGURATION of routing subsystem
DISABLE interface
ERROR counts
EVENT logging
INTERFACE statistics
LOG level
MEMORY statistics
NETWORK commands
PROTOCOL commands
QUEUE lengths
STATISTICS of network
TEST network
UPTIME of gateway
ELS-DISPLAY
DEBUG
PHDUMP

RANGE 2048-2111 +

```

Figure 213. GWCON Commands

6.2.2.1 GWCON Protocol Commands

A very common GWCON command is the PROTOCOL command. Figure 214, Figure 215 on page 264, and Figure 216 on page 264, illustrate the available commands for the IP, OSPF and BGP protocols.

```

RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol ip
RANGE 2048-2111 IP>?

RANGE 3746-port-number
ACCESS controls
CACHE
COUNTERS
DUMP routing tables
INTERFACE addresses
PACKET-FILTER
PING dest_addr [src_addr size ttl rate]
ROUTE given address
SIZES
STATIC routes
TRACEROUTE address
EXIT

RANGE 2048-2111 IP>exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *

```

Figure 214. GWCON IP Protocol Commands

```

RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol ospf
RANGE 2048-2111 OSPF>?

RANGE 3746-port-number
ADVERTISEMENT expansion
AREA summary
AS-EXTERNAL advertisements
DATABASE summary
DUMP routing tables
EXIT
INTERFACE summary
JOIN
LEAVE
MCACHE
MGROUPS
MSTATS
NEIGHBOR summary
PING dest_addr [src_addr size ttl rate]
ROUTERS
SIZE
STATISTICS
TRACEROUTE address
WEIGHT

RANGE 2048-2111 OSPF>exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *

```

Figure 215. GWCON OSPF Protocol Commands

```

RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol bgp
RANGE 2048-2111 BGP>?

RANGE 3746-port-number
DESTINATIONS
DUMP routing tables
NEIGHBORS
PATHS
PING dest_addr [src_addr size ttl rate]
SIZES
TRACEROUTE address
EXIT

RANGE 2048-2111 BGP>00367exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *

```

Figure 216. GWCON BGP Protocol Commands

6.2.3 CONFIG Commands

The Configuration Process (CONFIG) command level provides configuration facilities. It can also be used to list the active configuration and diagnose network problems. Config is accessed with the OPCON TALK 6 command.

Figure 217 on page 265 illustrates the CONFIG command level accessed from OPCON. Note that the command prompt becomes Config>.

```
RANGE 2048-2111 *talk 6
Gateway user configuration
Config>
```

Figure 217. CONFIG Command Level

The ? command lists the available CONFIG commands. Figure 218 gives an overview of the available CONFIG commands. Some of the commands we used in our testing are explained in 6.3, “TELNET Line Command Examples” on page 266.

```
Config>?
ADD
EVENT
EXIT
FEATURE
LIST
NETWORK
PATCH
PROTOCOL
SET
TIME of day params
UNPATCH
Config>
```

Figure 218. CONFIG Commands

6.2.3.1 CONFIG Protocol Commands

A very common CONFIG command is the PROTOCOL command. Figure 219, Figure 220 on page 266, and Figure 221 on page 266, illustrate the available commands after entering the IP, OSPF and BGP protocol level.

```
RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>?
LIST
CHANGE
DELETE
DISABLE
ENABLE
ADD
SET
MOVE
EXIT
IP config>exit
Config>exit
RANGE 2048-2111 *
```

Figure 219. CONFIG IP Protocol Commands

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ospf
Open SPF-Based Routing Protocol configuration console
OSPF Config>?
ADD
DELETE
DISABLE
ENABLE
EXIT
JOIN
LEAVE
LIST
SET
OSPF Config>exit
Config>exit

```

Figure 220. CONFIG OSPF Protocol Commands

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol bgp
Border gateway protocol user configuration
BGP Config>?
ADD
CHANGE
DELETE
DISABLE
ENABLE
LIST
MOVE
EXIT
BGP Config>exit
Config>exit

```

Figure 221. CONFIG BGP Protocol Commands

6.3 TELNET Line Command Examples

This section provides examples of a number of TELNET operator and configuration commands used during our testing.

The examples include:

- Testing IP connectivity using PING
- Displaying the CBSP routing table
- Displaying the 3746-9x0 processor IP routing cache
- A detailed display of a specific route
- Displaying interface (port) parameters
- Testing IP connectivity using Traceroute
- Displaying, adding, and deleting a static route
- Displaying, adding, and deleting an access control
- Displaying detailed OSPF information
- Displaying detailed frame relay information
- Displaying the 3746-9x0 IP router configuration

- Restarting the 3746-9x0 IP router

6.3.1 Using PING

Figure 222 illustrates how the PING command can be used to verify IP connectivity. Note that the IP PING command is continuously executed until the Enter key is pressed.

```
RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol ip
RANGE 2048-2111 IP>PING 192.168.157.9

PING 192.168.157.2 -> 192.168.157.9: 56 data bytes, ttl=64, every 1 sec.
56 data bytes from 192.168.157.9: icmp_seq=0. ttl=64. time=0. ms
56 data bytes from 192.168.157.9: icmp_seq=1. ttl=64. time=0. ms
56 data bytes from 192.168.157.9: icmp_seq=2. ttl=64. time=0. ms
56 data bytes from 192.168.157.9: icmp_seq=3. ttl=64. time=0. ms
56 data bytes from 192.168.157.9: icmp_seq=4. ttl=64. time=0. ms

[ENTER]
----192.168.157.9 PING Statistics----
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms

RANGE 2048-2111 IP>exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *
```

Figure 222. PING Command

6.3.2 Dumping the ARP Table

For problem determination it is often necessary to know the contents of the ARP table. Figure 223 shows the appropriate command.

```
RANGE 2112-2175 *talk 5

RANGE 2112-2175 +protocol arp

RANGE 2112-2175 ARP>dump
port/interface number or net name >2144
Protocol name or number to dump [IP]?[ENTER]
Hardware Address      IP Address      Refresh
40-00-37-46-21-44     9.24.104.168    2
40-00-52-00-51-39     9.24.104.171    1
10-00-5A-B1-B0-FE     9.24.104.12     5
10-00-5A-B1-AF-E9     9.24.104.109    5
00-04-AC-62-03-CE     9.24.104.50     5
10-00-5A-B1-AB-ED     9.24.104.119    5
00-04-AC-62-10-AC     9.24.104.188    5
10-00-5A-B1-4F-CB     9.24.104.124    5

RANGE 2112-2175 ARP>exit

RANGE 2112-2175 +exit

RANGE 2112-2175 *
```

Figure 223. ARP Table Dump

6.3.3 Dump CBSP Routing Table Display

Figure 224 displays the routing table on the CBSP.

```
RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol ip
RANGE 2048-2111 IP>dump
```

Type	Dest net	Mask	Cost	Age	Next hop(s)
Sbnt	9.0.0.0	FF000000	1	76963	None
Dir*	9.24.104.0	FFFFFFF0	1	77136	TKR/1
Dir*	10.0.0.0	FF000000	1	77135	TKR/3
Dir*	192.168.150.0	FFFFFFF0	1	77136	TKR/1
Dir*	192.168.152.0	FFFFFFF0	1	77163	TKR/0
Dir*	192.168.155.0	FFFFFFF0	1	76197	FR/6
Sbnt	192.168.156.0	FFFFFFF0	1	77103	None
SPF	192.168.156.3	FFFFFFF0	1	77123	192.168.157.9 (2)
SPF*	192.168.157.0	FFFFFFF0	1	77135	TKR/2
Dir*	192.168.202.0	FFFFFFF0	1	77139	ESCON/4

```
Routing table size: 768 nets (46080 bytes), 10 nets known
RANGE 2048-2111 IP>exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *
```

Figure 224. GWCON CBSP Routing Table Display

Note: The DUMP command can only be executed on the CBSP. Make sure your command is directed to the appropriate processor (CBSP) (see the discussion in 5.10.2, “TELNET Command Line Interface” on page 218). To display the routing caches maintained on the other processors use the CACHE command (see 6.3.4, “3746-9x0 Processor IP Routing Cache Display”).

6.3.4 3746-9x0 Processor IP Routing Cache Display

Figure 225 on page 269 displays the IP routing cache for the 3746-9x0 processor that controls token-ring interface 2176. Note how the RANGE command was used to redirect the CACHE to the proper processor.


```

RANGE 2048-2111 *range 2176
RANGE 2176-2239 *talk 5
RANGE 2176-2239 +protocol ip
RANGE 2176-2239 IP>cache

Destination          Usage      Next hop
9.24.104.171         22595    9.24.104.171    (Pro/0)
192.168.152.1         303      192.168.152.1    (Pro/0)
192.168.152.2         304      192.168.152.2    (Pro/0)
10.0.0.8              303      10.0.0.8         (TKR/3)
192.168.202.9         606      192.168.202.9    (Pro/0)
192.168.202.19        606      192.168.202.19   (Pro/0)
192.168.156.3         304      192.168.156.3    (Pro/0)
192.168.155.1         302      192.168.155.1    (Pro/0)

RANGE 2176-2239 IP>exit
RANGE 2176-2239 +exit
RANGE 2176-2239 *range 2080
RANGE 2048-2111 *

```

Figure 225. IP Routing Cache Display

Note: The DUMP command can be used to display the routing table on the CBSP (see 6.3.3, “Dump CBSP Routing Table Display” on page 268).

6.3.5 Specific Route Display

Figure 226 displays detailed routing information for IP address 192.168.156.9.

```

RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol ip
RANGE 2048-2111 IP>route 192.168.155.9

Destination: 192.168.155.0
Mask: 255.255.255.0
Route type: Dir
Distance: 1
Age: 77326
Tag: 0
Next hop(s): 192.168.155.9    (FR/6)

RANGE 2048-2111 IP>exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *

```

Figure 226. IP Detailed Route Display

6.3.6 IP Protocol Display

Figure 227 on page 270, and Figure 228 on page 270 illustrates the IP address, the CCM port names and subnet mask assigned to the 3746 IP ports. Similar information can be retrieved from GWCON and CONFIG.

```

RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol ip
RANGE 2048-2111 IP>interface

Interface          IP Address(es)    Mask(s)
TKR/0-      2080 192.168.152.3    255.255.255.0
TKR/1-      2144 9.24.104.168      255.255.255.0
              192.168.150.3    255.255.255.0
TKR/2-      2176 192.168.157.2    255.255.255.0
TKR/3-      2208 10.0.0.3        255.0.0.0
ESCON/4-    2240 192.168.202.8    255.255.255.0
PPP/5-      2310 0.0.0.0          0.0.0.0
FR/6-      2311 192.168.155.9    255.255.255.0

RANGE 2048-2111 IP>exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *

```

Figure 227. GWCON IP Protocol Interface Display

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>list addresses

0 PORT2080 192.168.152.3    255.255.255.0    Local wire broadcast, fill
1 PORT2144 192.168.150.3    255.255.255.0    Local wire broadcast, fill
              9.24.104.168    255.255.255.0    Local wire broadcast, fill
2 PORT2176 192.168.157.2    255.255.255.0    Local wire broadcast, fill
3 PORT2208 10.0.0.3        255.0.0.0        Local wire broadcast, fill
4 P2240IP 192.168.202.8    255.255.255.0    Local wire broadcast, fill
5 PP2310IP 0.0.0.5          0.0.0.0          Local wire broadcast, fill
6 PP2311IP 192.168.155.9    255.255.255.0    Local wire broadcast, fill

IP config>exit
Config>exit

RANGE 2048-2111 *

```

Figure 228. CONFIG IP Protocol List Addresses Display

6.3.7 Using Traceroute

The Traceroute command is very helpful in displaying the network route taken for a specific IP address. Figure 229 on page 271 illustrates a route trace.

```

RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol ospf
Open SPF-Based Routing Protocol console

RANGE 2048-2111 OSPF>traceroute 192.168.156.9
TRACEROUTE 192.168.156.9: 56 data bytes
 1 192.168.152.3 0 ms !N 0 ms !N 100 ms !N

RANGE 2048-2111 OSPF>exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *

```

Figure 229. Traceroute Display

6.3.8 Static Routes

Figure 230, Figure 231 on page 272, and Figure 232 on page 272 illustrates how static routes can be added, displayed, and deleted using CONFIG IP Protocol commands.

Note

No 3746 IP router restart is required to activate the Add Route and Delete Route commands, as these commands are executed immediately.

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>add route
IP destination []? 192.168.156.0
Address mask [255.255.255.0]? 255.255.255.0
Via gateway at []? 192.168.157.9
Cost [1]? 2
Via gateway 2 at []? [ENTER]
IP config>exit
Config>exit
Config: writing block in sram file: type=2, length=360
Config: writing block in sram file: type=3, length=2E0
Config: writing block in sram file: type=5, length=700
Config: writing block in sram file: type=D, length=40
Config: writing block in sram file: type=0, length=0

RANGE 2048-2111 *

```

Figure 230. Adding a Static Route

```

RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol ip
RANGE 2048-2111 IP>static

Net          Mask          Cost  Next hop
192.168.156.0 255.255.255.0 2      192.168.157.9

RANGE 2048-2111 IP>exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *

```

Figure 231. Displaying a Static Route

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>delete route
IP destination []? 192.168.156.0
Address mask [255.255.255.255]? 255.255.255.0
IP config>exit
Config>exit
Config: writing block in sram file: type=2, length=360
Config: writing block in sram file: type=3, length=2E0
Config: writing block in sram file: type=5, length=6E0
Config: writing block in sram file: type=D, length=40
Config: writing block in sram file: type=0, length=0

RANGE 2048-2111 *

```

Figure 232. Deleting a Static Route

6.3.9 Using Access Controls

The definition of IP access controls enables you to control IP forwarding in your network. For an explanation of how access controls work, refer to 5.12, “3746-9x0 Security” on page 229. Figure 233 on page 273, Figure 234 on page 274, Figure 235 on page 274, Figure 236 on page 275, and Figure 237 on page 275 illustrate how you can prevent IP traffic from the source IP address 192.168.157.7 host within subnet 192.168.152.0 being forwarded by the 3746 IP router. This is independent of the protocol and TCP/UDP port being used. The figures show how to add an access control entry, display the access control list, and delete an access control entry.

Important

- New entries are always added at the bottom of the access control list. Be aware of the access control entry that enables any-to-any traffic (source/destination 0.0.0.0). Make sure that this entry is always present and moved to the bottom of the access list before activating your access list. Unless you move this default entry, your changes will have no effect. This is because the default entry will match every packet and cause all packets to be forwarded, your new entry (which is lower down in the list than the default entry) will never be compared against an IP packet.
- The CONFIG command level must be exited and the 3746 IP router restarted to activate the Add Access-control and Delete Access-control commands. For a discussion of the OPCON RESTART command see 6.3.14, “3746-9x0 IP Router Restart” on page 280.

```
RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>list access-controls
Access Control is: enabled
List of access control records:

   Ty Source      Mask      Destination      Mask      Beg End   Beg End
1  I  0.0.0.0      00000000  0.0.0.0          00000000  0  255  0  65535
IP config>exit
Config>exit
RANGE 2048-2111 *
```

Figure 233. Display an Access Control

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>add access-control
Enter type [E]? E
Internet source [0.0.0.0]? 192.168.150.7
Source mask [255.255.255.255]? 255.255.255.255
Internet destination [0.0.0.0]? 192.168.152.0
Destination mask [255.255.255.255]? 255.255.255.0
Enter starting protocol number ([CR] for all) [-1]? [ENTER]
IP config>
IP config>list access-controls
Access Control is: enabled
List of access control records:

      Ty Source      Mask      Destination      Mask      Beg End      Beg End
1  I  0.0.0.0      00000000  0.0.0.0      00000000  0  255  0  65535
2  E  192.168.150.7 FFFFFFFF  192.168.152.0 FFFFFF00  0  255  0  65535
IP config>move access-control
Enter index of control to move [1]1
Move record AFTER record number [0]? 2
About to move:

      Ty Source      Mask      Destination      Mask      Beg End      Beg
1  I  0.0.0.0      00000000  0.0.0.0      00000000  0  255  0  65535
to be after:
2  E  192.168.150.7 FFFFFFFF  192.168.152.0 FFFFFF00  0  255  0  65535
Are you sure this is what you want to do(Yes or [No]): yes
IP config>exit
Config>exit
Config: writing block in sram file: type=2, length=360
Config: writing block in sram file: type=3, length=2E0
Config: writing block in sram file: type=5, length=760
Config: writing block in sram file: type=D, length=40
Config: writing block in sram file: type=0, length=0
RANGE 2048-2111 *

```

Figure 234. Adding, Displaying, Moving, and Activating an Access Control

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>list access-controls

Access Control is: enabled
List of access control records:

      Ty Source      Mask      Destination      Mask      Beg End      Beg End
1  E  192.168.150.7 FFFFFFFF  192.168.152.0 FFFFFF00  0  255  0  65535
2  I  0.0.0.0      00000000  0.0.0.0      00000000  0  255  0  65535

IP config>exit
Config>exit
RANGE 2048-2111 *

```

Figure 235. Display an Access Control

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>delete access-control

Enter index of access control to be deleted [1]? 1

      Ty Source      Mask      Destination      Mask      Beg End      Beg End
      1 E 192.168.150.7 FFFFFFFF 192.168.152.0 FFFFFF00 0 255 0 65535
Are you sure this record you want deleted(Yes or [No]): yes
Deleted
IP config>exit
Config>exit
CONFIG: writing block in sram file: type=2, length=360
CONFIG: writing block in sram file: type=3, length=420
CONFIG: writing block in sram file: type=5, length=6C0
CONFIG: writing block in sram file: type=C, length=100
CONFIG: writing block in sram file: type=D, length=40
CONFIG: writing block in sram file: type=0, length=0
RANGE 2048-2111 *

```

Figure 236. Deleting an Access Control

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>set access-control on
IP config>exit
Config>exit
CONFIG: writing block in sram file: type=2, length=360
CONFIG: writing block in sram file: type=3, length=460
CONFIG: writing block in sram file: type=5, length=760
CONFIG: writing block in sram file: type=C, length=120
CONFIG: writing block in sram file: type=D, length=40
CONFIG: writing block in sram file: type=F, length=80
CONFIG: writing block in sram file: type=0, length=0

```

Figure 237. Enable an Access Control

6.3.10 Using Filters

For a detailed description of IP filters refer to 5.12, “3746-9x0 Security” on page 229. No gateway restart is required to activate filters. A filter is specified as an address (the destination IP address) and a subnet mask, and if a packet matches a filter then it is discarded. In the following example we add a filter that prevents packets for the destination address 192.168.157.9 from passing through the router.

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>add filter
IP destination []? 192.168.157.9
Address mask [255.255.255.0]? 255.255.255.255
IP config>exit
Config>exit
CONFIG: writing block in sram file: type=2, length=360
CONFIG: writing block in sram file: type=3, length=460
CONFIG: writing block in sram file: type=5, length=740
CONFIG: writing block in sram file: type=C, length=120
CONFIG: writing block in sram file: type=D, length=40
CONFIG: writing block in sram file: type=F, length=80
CONFIG: writing block in sram file: type=0, length=0

RANGE 2048-2111 *

```

Figure 238. Adding a Filter

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>list filter

Filter address 192.168.157.9, 255.255.255.255

IP config>exit
Config>exit
RANGE 2048-2111 *

```

Figure 239. Displaying Defined Filters


```

RANGE 2048-2111 *talk 5
RANGE 2048-2111 +protocol ip
RANGE 2048-2111 IP>dump

```

Type	Dest net	Mask	Cost	Age	Next hop(s)
Sbnt	9.0.0.0	FF000000	1	5166	None
Dir*	9.24.104.0	FFFFFF00	1	5173	TKR/1
Dir*	10.0.0.0	FF000000	1	5172	TKR/3
Dir*	192.168.150.0	FFFFFF00	1	5173	TKR/1
Dir*	192.168.152.0	FFFFFF00	1	5196	TKR/0
Dir*	192.168.155.0	FFFFFF00	1	5166	FR/6
Sbnt	192.168.156.0	FFFFFF00	1	5136	None
SPF	192.168.156.3	FFFFFFF	1	5160	PPP/5
SPF*	192.168.157.0	FFFFFF00	1	5172	TKR/2
Fltr	192.168.157.9	FFFFFFF	16	390	SINK/0
Dir*	192.168.202.0	FFFFFF00	1	5175	ESCON/4

Routing table size: 768 nets (46080 bytes), 11 nets known

```

RANGE 2048-2111 IP>exit
RANGE 2048-2111 +exit
RANGE 2048-2111 *

```

Figure 240. Is the Filter in the Routing Table Too?

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>delete filter
IP destination []? 192.168.157.9
Address mask [255.255.255.0]? 255.255.255.255
IP config>exit
Config>exit
CONFIG: writing block in sram file: type=2, length=360
CONFIG: writing block in sram file: type=3, length=460
CONFIG: writing block in sram file: type=5, length=720
CONFIG: writing block in sram file: type=C, length=120
CONFIG: writing block in sram file: type=D, length=40
CONFIG: writing block in sram file: type=F, length=80
CONFIG: writing block in sram file: type=0, length=0
RANGE 2048-2111 *

```

Figure 241. Deleting a Filter

6.3.11 Protocol Global Configuration Display

Within a protocol process level in CONFIG, it may be desirable to list the current router configuration. Figure 242 on page 278 gives a detailed overview of the OSPF protocol configuration.

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ospf
Open SPF-Based Routing Protocol configuration console
OSPF Config>list all

--Global configuration--
OSPF Protocol:      Enabled
# AS ext. routes:   100
Estimated # routers: 50
External comparison: Type 2
AS boundary capability: Enabled
Import external routes: BGP RIP STA DIR SUB
                    BGP auto-tag generation disabled
Orig. default route: No (0,0.0.0.0)
Default route cost:  (1, Type 2)
Default forward. addr.: 0.0.0.0
Multicast forwarding: Disabled

--Area configuration--
Area ID      AuType      Stub? Default-cost Import-summaries?
0.0.0.0      0=None          No           N/A             N/A

--Interface configuration--
IP address    Area      Cost Rtrns TrnsDly Pri Hello Dead
192.168.157.3 0.0.0.0    1     5     1     1    10    40

OSPF Config>exit
Config>exit
RANGE 2048-2111 *

```

Figure 242. OSPF Protocol Configuration Display

6.3.12 Display Frame Relay Information

To display the PVCs use Figure 243.

```

RANGE 2048-2111 *talk 5

RANGE 2048-2111 +network 2311
Frame Relay Console

RANGE 2304-2367 FR>list permanent-virtual-circuit

  Circuit      Frames      bytes
  Nb      State      Sent  Received      Sent  Received
  -----
    32      Active      184      3      15000      90
    33      Active      16      0      1324      0

RANGE 2304-2367 FR>

```

Figure 243. Display PVCs

Figure 244 on page 279 shows how to display the IARP table for a frame relay port.

```

RANGE 2304-2367 *talk 5

RANGE 2304-2367 +network 2311
Frame Relay Console

RANGE 2304-2367 FR>list iarp
  ARP Table:  DLCI #  <>  IP address
                32      192.168.155.1
                33      192.168.155.33

RANGE 2304-2367 FR>

```

Figure 244. Displaying the IARP Table for a Frame Relay Port

Figure 245 shows how to display all user-defined address pairs.

```

RANGE 2304-2367 *talk 6
Gateway user configuration
Config>network 2311
Frame Relay user configuration
FR Config>list protocol-address

                        Frame Relay Protocol Address Translations

Protocol Type          Protocol Address          Circuit Number
-----
      IP                192.168.155.33                33
FR Config>

```

Figure 245. Displaying All Configured Address Pairs

6.3.13 Display 3746 IP Router Configuration

Figure 246, and Figure 247 on page 280 display the devices and the protocol configuration of the 3746 IP router while in the CONFIG command level.

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>list devices

Ifc 0 (Token Ring): PORT 2080, LIM 2, SLOT 32, NAME P2080IP
Ifc 1 (Token Ring): PORT 2144, LIM 3, SLOT 32, NAME P2144IP
Ifc 2 (Token Ring): PORT 2176, LIM 4, SLOT 0, NAME P2176IP
Ifc 3 (Token Ring): PORT 2208, LIM 4, SLOT 32, NAME P2208IP
Ifc 4 (Escon): PORT 2240, LIM 5, SLOT 0, NAME P2240IP

Config>exit
RANGE 2048-2111 *

```

Figure 246. CONFIG List Devices

```

RANGE 2048-2111 *talk 6
Gateway user configuration
Config>list configuration
Initial logging level: 11

Configurable Protocols:
Num Name Protocol
0 IP DOD-IP
3 ARP Address Resolution
10 BGP Border Gateway Protocol
12 OSPF Open SPF-Based Routing Protocol

795520 bytes of configuration memory free

Config>exit
RANGE 2048-2111 *

```

Figure 247. CONFIG List Configuration

Figure 248 displays the device and protocol configuration using GWCON commands.

```

RANGE 2048-2111 *talk 5
RANGE 2048-2111 +configuration
Portable I80386 C Gateway USIBMRA.NN061A
V15.1[]

Num Name Protocol
0 IP DOD-IP
3 ARP Address Resolution
10 BGP Border Gateway Protocol
12 OSPF Open SPF-Based Routing Protocol

5 Networks:
Net Interface MAC/Data-Link Hardware State
0 TKR/00- 2080 Token-Ring/802.5 ProNET-4/16 Up
1 TKR/01- 2144 Token-Ring/802.5 ProNET-4/16 Up
2 TKR/02- 2176 Token-Ring/802.5 ProNET-4/16 Up
3 TKR/03- 2208 Token-Ring/802.5 ProNET-4/16 Up
4 ESCON/04- 2240 ESCON ESCON Up

RANGE 2048-2111 +exit
RANGE 2048-2111 *

```

Figure 248. GWCON Configuration

6.3.14 3746-9x0 IP Router Restart

The 3746 IP router is restarted using the OPCON RESTART command illustrated in Figure 249 on page 281.

```
RANGE 2048-2111 *restart
Are you sure you want to restart the gateway? (Yes or [No]):
RANGE 2048-2111 *yes
Restart of gateway in progress...
RANGE 2048-2111 *
```

Figure 249. 3746-9x0 IP Router Restart

A restart is, for example, required to activate the CONFIG commands you have entered. Note that adding/deleting static routes, IP filters, and OSPF neighbors does not require a router restart, as these commands are executed immediately.

Chapter 7. Using CCM - IP Configuration Examples

The following sections detail a number of connection scenarios that can be used to integrate the 3746 IP router within your IP network, and connect IP hosts and routers to the 3746 IP router. The required parameter settings for the 3746 IP router have also been included.

The chapter is divided in three major topics. At the beginning there is the configuration of the service ring itself. Due to the IP protocol being used for communication between NNP, CBSP (Port 2080) and SP it is necessary to configure this part of the machine first. Usually the machine comes up with default addresses for all its components connected to the service LAN. But if you want to connect service LANs together, you must ensure that the addresses in each LAN are unique. The second section shows basic configuration, this means supplying each adapter with the necessary parameters (for example, ring speed or IP address). The last topic describes all advanced configuration, for example, routing.




Figure 250. 3746 IP Implementation Test Network

Figure 250 illustrates our test environment. It consists of a 3746-900 IP router with four token-ring interfaces and one ESCON interface. The 3746-900 only performs IP functions, and no APPN functions have been configured. Each of the token-ring interfaces is identified by the dynamic routing protocol being used on it:

- Service LAN (token-ring port 2080)

The service LAN provides connectivity to the SP and NNP. A single IP address, 192.168.152.3 with subnet mask 255.255.255.0, has been defined on this port. See Figure 251 on page 285.

Important

It is mandatory for proper 3746 IP operation that port 2080 is enabled for IP transport. The SP, NNP, and port 2080 must all be part of the same subnet.

- OSPF LAN (token-ring port 2176)

Token-ring port 2176 is used to connect to the OSPF LAN. To enable the exchange of OSPF dynamic routing information on this LAN, OSPF has been enabled on this port. A single IP address, 192.168.157.2 with subnet mask 255.255.255.0, has been defined on this port. See Figure 276 on page 297.

- BGP LAN (token-ring port 2208)

Token-ring port 2208 is used to connect to the BGP LAN. To enable the exchange of BGP routing information on this LAN, BGP has been enabled for this port. A single IP address, 10.0.0.3 with subnet mask 255.0.0.0, has been defined on this port. See Figure 281 on page 300.

- Office LAN (token-ring port 2144)

Token-ring port 2144 connects to the ITSO Raleigh office LAN backbone. Two IP addresses have been defined on this port:

- 9.24.104.168 with subnet mask 255.255.255.0

Being part of the IBM internal 9.0.0.0 production network enabled us to access the 3746-9x0 SNMP functions from a NetView/6000 management station. No dynamic routing protocol has been associated with this IP address to prevent IP routing information from our test network interfering with the production 9.0.0.0 network.

- 192.168.150.3 with subnet mask 255.255.255.0

The definition of address 192.168.150.3 allows access to the 3746 IP router from office LAN-attached workstations defined within subnet 192.168.150. RIP functions have been enabled for this address. See Figure 269 on page 293.

- ESCON Connection (ESCON port 2240)

On the ESCON connection we have assigned IP address 192.168.202.8, using subnet mask 255.255.255.0, to the 3746 (see Figure 286 on page 303). To enable the exchange of dynamic routing information we have activated RIP on this interface. Note that native IP over ESCON requires TCP/IP for MVS (≥V3R1) or OS/390 (≥V2R4). RIP is the only dynamic routing protocol supported by TCP/IP for MVS or the OS/390 TCP/IP stack.

Important

Part of our test environment, for example, the use of three different routing protocols, is unlikely to occur in a production environment. The purpose of this test environment was to enable us to show the various configuration steps required, verify the interaction of the various routing protocols, and do some basic connectivity testing.

7.1 Defining the Service LAN

Token-ring port 2080 is used by the 3746-9x0 to connect to the service LAN. Figure 251 shows the configuration of the service LAN.




Figure 251. Service LAN Configuration

In our test scenario, we have isolated the service LAN from the rest of our network. We have no other connections into our service LAN. This allows us to control access to the service LAN, and therefore we can be certain that no other traffic will create a load on the service LAN. In some cases (TELNET access to the NNP for example) we may need IP connectivity to the service LAN, which is provided via the 3746. From the workstation which we wish to use to access the NNP, we must define a route to network 192.168.152.0 by specifying an IP address on the 3746 as the router for that network. Figure 252 shows an example of the definitions needed on a workstation connected via port 2144 to the 3746.

```
D:\> route add net 192.168.152.0 9.24.104.168 1
add net 192.168.152.0: router 9.24.104.168

D:\> netstat -r
```

destination	router	refcnt	use	flags	snmp metric	intrf
192.168.152.0	9.24.104.168	0	99	U	-1	lan0

Figure 252. Workstation Definitions

Therefore, access to the service LAN is always through the 3746. The NNP automatically has a default route which points to port 2080, this provides a return route back to the 3746. There is no reason why further 3746 machines should not also be provided routed IP access to the service LAN of a different 3746. This allows access to the SP and NNP even when their 3746 is unavailable or unreachable. An example of this is shown in Figure 253 on page 286.




Figure 253. Alternate Access to the Service LAN

As in our configuration we wanted to use an IP subnet that was different from the default defined during box installation. We also had to change the SP and NNP IP addresses (see 5.10, “3746-9x0 IP Management” on page 214). Fortunately all these addresses are placed at the same panel in the service processor menu so they can be changed easily.




Figure 254. Accessing Service Processor Function Menu

Figure 254 shows how to obtain the service processor menu. After the menu appears select **Configuration Management** (see Figure 255).




Figure 255. Service Processor Menu

In the configuration management subtree select **SP Customization** (see Figure 256).




Figure 256. SP Customization

At this time, we only want to change the Service LAN Addresses, therefore select only this check box. Deselect all others if necessary, then click **Next** (see Figure 257 on page 288).




Figure 257. Customization Selection Menu

This brings you to the panel shown in Figure 258, enter the new addresses you wish to use. In the SP default router field you can define the required default router for the SP. This is important if you wish to access the SP from outside of the service LAN, for example using DCAF, and it allows the SP to route the IP packets back to their source. In our example we have used port 2080 as the default router. The NNP also needs a default router for the cases where it is accessed via TELNET from outside the service LAN. The default route is added to the NNP automatically and it always points to port 2080. It is not user configurable. After all changes are done select **Next** (see Figure 258).

	IP address	Subnet mask	Hostname	UAA/LAA
Service Processor:	192.168.152.1	255.255.255.0	SP00024	40003746F9F9
NNP- A:	192.168.152.2	255.255.255.0	CA191805	
NNP- B:	not installed			
TIC3 2080:	192.168.152.3	255.255.255.0		
SP default router:	192.168.152.3			
MAE:				
LAN Manager				
Do you have a LAN manager?			C&SM LAN ID: MOSSE	
☐ Yes ☒ No				

<<Previous Next>> Help

Figure 258. Service LAN Address

Again the Service Processor Customization Selection menu appears. Now select **Close** (see Figure 259 on page 289).




Figure 259. Customization selection menu

A warning appears. To accept your changes click **Yes** (see Figure 260).

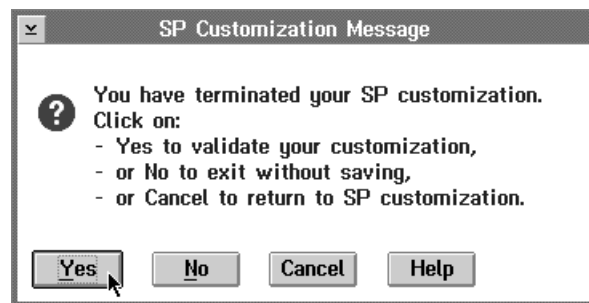


Figure 260. SP Customization Message

Now the changes you made are applied. While this is happening, the window shown in Figure 261 is displayed.




Figure 261. Service Processor Customization - Progress Window

If no problems occur, the window in Figure 262 is displayed. Click the **OK** button.




Figure 262. Service Processor Customization Competition

The definition of the IP addresses of the service LAN is now complete. If you want to change any other port 2080 parameters, use the procedure described in 7.2, “Defining Token-Ring Interface Port 2080” on page 290.

7.2 Defining Token-Ring Interface Port 2080

Token-ring port 2080 is used by the 3746-9x0 to connect to the service LAN. To configure this port use the following procedure. Remember that the IP address configuration is done in the Service processor menu.

For all ports, there are two methods available to start configuration:

1. Select the port ICON, then select **Configuration->Coupler**.
2. Double click on the port Icon.

For PORT 2080, an additional step is necessary. When using APPN, definitions need to be made for an NPALU station and the service processor station. These are not needed when only running IP. Figure 263 shows how to access the dialog to define these APPN stations on PORT 2080. Select the **Stations on coupler/port 2080** option in the Configuration pull down menu.




Figure 263. Port 2080 Configuration Option

If you double click the PORT 2080 icon as shown in Figure 264 on page 291, then the dialog shown in Figure 265 on page 291 appears to allow configuration of the APPN stations.




Figure 264. Coupler 2080

The Service Processor option should be checked if APPN connectivity is required to the SP, for example, when DCAF over SNA is used to the SP, or if any APPN reporting of alerts is desired. CCM automatically creates an APPN station definition for the SP if this option is enabled. The NPM option should be checked if an NPALU station is to be defined on your behalf by CCM.

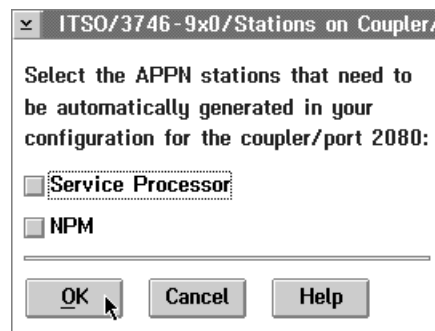


Figure 265. Stations on Coupler/Port 2080

After clicking on OK, the CCM screen shown in Figure 266 on page 292 is displayed.




Figure 266. Token-Ring Port 2080 Configuration

We enabled IP and set the IP port name to PORT2080. The IP name is a mandatory field when using a port for IP. The maximum transmission unit (MTU) size range has been defined as 2052. The range of MTU sizes available depends on the token-ring speed setting.




Figure 267. IP over Token-Ring Parameters

Figure 267 illustrates that the IP address for this port can not be changed. This address is defined during the SP configuration process (see 7.1, “Defining the Service LAN” on page 285).




Figure 268. Coupler/Port 2080 Configured

After entering the IP parameters and selecting **OK** we return to the primary CCM configuration panel displayed in Figure 268. Note that on the coupler icon for port 2080 a small check mark has appeared.

7.3 Defining Token-Ring Interface Port 2144

The test environment illustrated in Figure 269 uses token-ring port 2144 to attach to the 16-Mbps ITS0 office LAN.




Figure 269. Office LAN Configuration

To enable the use of RIP on this port without interfering with the production network, two IP addresses have been defined. One address is part of a test

environment (IP address 192.168.150.3), and the other is part of the production network (IP address 9.24.104.168). The following screens show the CCM configuration steps.




Figure 270. Port 2144 Configuration Option

After initiating CCM configuration from the 3746-9x0 menu (see 6.1.4, "Becoming Familiar with the CCM User Interface" on page 244), coupler port 2144 can be selected simply by double-clicking on address 2144 as shown in Figure 270.




Figure 271. 3745/3746 Parameters

When configuring the first port, other than port 2080, CCM will require you to enter the 3746-9x0 model and the 3745 operation (see Figure 271). The 3745 parameters only apply for a 3746 model 900. As shown, our test environment uses a 3746 Model 900 to which a 3745 Model 41A running in twin dual mode has been attached.



Figure 272. Coupler Type

Once the 3746 model and, if applicable, the 3745 parameters have been specified CCM will require you to identify the coupler type (see Figure 272). Because we have a dual-CCU 3745 attached to our 3746-900, port 2112 is reserved for a CBC and thus cannot be configured.




Figure 273. Token-Ring Port 2144 Configuration

Figure 273 illustrates our definitions for port 2144. We have disabled APPN, enabled IP, and specified IP name PORT2144. Port 2144 connects to a 16-Mbps LAN using token-ring MAC address 400037462144. The MTU size has been set to 4060. Additional IP parameters can be configured by selecting IP parameters.

ITS0/3746-9x0/IP Over Token-Ring Parameters	
Port: 2144	Name: PORT2144
General	
Automatic reactivation? ☒ Yes ☐ No	
☒ Enable source routing	
RIF timer: 120 seconds [0-4096]	
Addresses	
IP address:	Subnet mask:
192.168.150.3	255.255.255.0
192.168.150.3	255.255.255.0
9.24.104.168	255.255.255.0
Add Modify Delete	
OK Save as defaults Cancel Help	

Figure 274. IP over Token-Ring Parameters

Selecting IP parameters enables you to set the IP address along with the aging timer for the routing information (RIF) cache. The *enable source routing* check box determines whether frames are originated with a token-ring routing information field (RIF). The RIF is required if a frame is to pass source route (SR) bridges. However, some stations cannot handle frames with RIF.

Figure 274 illustrates the two IP addresses and subnet masks. All other parameters remain at their default settings. Return to the main CCM configuration screen (Figure 275 on page 297) by selecting **OK**. Note that after being configured the icon for coupler 2144 contains a small check mark.




Figure 275. Coupler/Port 2144 Configured

7.4 Defining Token-Ring Interface Port 2176

The test environment illustrated in Figure 276 shows token-ring port 2176 which is attached to the OSPF LAN.




Figure 276. OSPF LAN Configuration

The following screens display the CCM configuration steps.




Figure 277. Coupler/Port 2176

After initiating CCM configuration from the 3746-9x0 menu (see 6.1.4, “Becoming Familiar with the CCM User Interface” on page 244), coupler port 2176 can be selected by double-clicking on the icon for coupler 2176 (see Figure 277).




Figure 278. Token-Ring Port 2176 Configuration

Figure 278 illustrates our definitions for port 2176. We disabled APPN, enabled IP, and assigned IP name PORT2176. Port 2176 connects to a 16-Mbps LAN with MAC address 400037462176. The MTU size is set to 4060, which is consistent with the definitions in the rest of the network. Further IP parameters can be configured by selecting IP parameters.

ITS0/3746-9x0/IP Over Token-Ring Parameters

Port: 2176 Name: PORT2176

General

Automatic reactivation? ☒ Yes ☐ No

☒ Enable source routing

RIF timer: 120 seconds [0-4096]

Addresses

IP address: 192.168.157.2 Subnet mask: 255.255.255.0

192.168.157.2 255.255.255.0

Buttons: Add, Modify, Delete, OK, Save as defaults, Cancel, Help

Figure 279. IP over Token-Ring Parameters

Selecting IP parameters enables you to set the IP address and the aging timer value for the Routing Information Timer (RIF) cache. Figure 279 depicts the definition of IP address 192.168.150.3, subnet mask 255.255.255.0, and a default RIF timer value.

ITS0/3746-9x0/CCM

File Configuration Management Options Help

Active configuration is: SG24-4845-01 (Nov-24-1997)

Opened configuration is: SG24-4845-01 (Nov-24-1997)

Opened reference CDF-E is:none

2752	2784	2816	2848	2880	2912	2944	2976	3008	3040	3072	3104
2368	2400	2432	2464	2496	2528	2560	2592	2624	2656	2688	2720
2048	2080 TIC3	2112	2144 TIC3	2176 TIC3	2208 TIC3	2240	2272	2304	2336		

Coupler 2176

Figure 280. Coupler/Port 2176 Configured

After selecting **OK** one returns to the main CCM configuration menu. Note that after configuration of port 2176, the icon for coupler 2176 contains a small check mark.

7.5 Defining Token-Ring Interface Port 2208

The test environment illustrated in Figure 281 contains port 2208 connected to a 16-Mbps LAN that is used for the exchange of BGP routing information (see 7.10, “BGP Parameter Configuration” on page 321).




Figure 281. BGP LAN Configuration

The following screens show the CCM configuration steps.




Figure 282. Coupler/Port 2208

After initiating CCM configuration from the 3746-9x0 menu (see 6.1.4, “Becoming Familiar with the CCM User Interface” on page 244) coupler port 2208 can be selected by double-clicking (see Figure 282) on the icon.

Figure 283. Token-Ring Port 2208 Configuration

Figure 283 shows the definitions for port 2208. We disabled APPN, enabled IP, and defined IP name PORT2208. Port 2208 connects to a 16-Mbps LAN using MAC address 400037462208. The MTU size has been set to 4060, as we did on all token-ring ports. Additional IP parameters can be set by selecting IP parameters.

IP address:	Subnet mask:
10.0.0.3	255.0.0.0

Figure 284. IP over Token-Ring Parameters

Selecting IP parameters enables us to set the IP address, the Subnet mask, and a value for the RIF cache aging timer. Figure 284 illustrates the use of IP address 10.0.0.3, subnet mask 255.0.0.0, and a default RIF timer value.




Figure 285. Coupler/Port 2208 Configured

After selecting OK the main CCM configuration screen reappears. Note that port 2208 now contains a small check mark.

7.6 Defining ESCON Interface Port 2240

The test environment illustrated in Figure 250 on page 283 uses native IP over ESCON on port 2240. A single IP link station has been defined to connect to TCP/IP for MVS. Between the host and 3746-900 is an ESCON director. The ESCON line will have RIP enabled. For more information about the ESCON definition see 5.4.5, "ESCON - IP Configuration Example" on page 168.




Figure 286. ESCON Configuration

The following screens display the CCM configuration.





Figure 287. Coupler/Port 2240

After initiating CCM configuration from the 3746-9x0 menu (see 6.1.4, “Becoming Familiar with the CCM User Interface” on page 244), coupler port 2240 can be selected by double-clicking the address as shown in Figure 287.

ITSO/3746-9x0/ESCON Port Configuration

Example: Single ESCD



Configure ESCON Port 2240

Network: ☐ APPN ☒ IP ☐ SNA/Subarea

Fiber status? ☒ Enable ☐ Transmit OLS ☐ Disable

Port name:

Automatic reactivation? ☐ Yes ☒ No

Number of host links: 1 NPA eligible? ☐ Yes ☒ No

Comments (optional):

ESCON Director (ESCD)

Port attached to an ESCD? ☒ Single ☐ Chained ☐ None

ESCD number: hex ESCD model:

Control Unit Link Address (LINK): hexadecimal

OK Host links... Delete port Cancel Help

Figure 288. ESCON Port 2240 Configuration

Figure 288 illustrates our definitions for the ESCON port. We disabled APPN, enabled IP, and assigned the IP name P2240IP. We select Port attached to a single ESCON director, with ESCON director number E1, ESCON director model 9032 and the Control Unit Link address EA. Then select the **Host links** option to define an ESCON host link (and then on the host link, a link station). This displays the dialog shown in Figure 289.


ITSO/3746-9x0/ ESCON Host Links Configuration - Port 2240

Port name

APPN:

IP: P2240IP

Example: Basic Mode



Number of host links: 1 Control Unit Link Address: EA

Configure a Host Link

Network? ☐ APPN (A) ☒ IP (I) ☐ SNA/Subarea (S)

Host link name:

Host mode? ☒ Basic ☐ LPAR ☐ EMIF

Host name: CHPID: hex

Partition name: Partition number: ☐ Dynamic ☐ Defined hex

Host Link Address (HLA): ☐ Dynamic ☒ Defined hex

Add Modify

Host Links Already Configured

No.	Network	APPN name	IP name	Host mode
1	IP	N/A	HL2240IP	Basic

Delete APPN parameters... Stations...

OK Cancel Help

Figure 289. ESCON Host Links Configuration Panel - Port 2240


We defined IP link HL2240IP to SYS6 host (running in basic mode) using CHPID 19. Host Link Address is E2 and then selected *stations* to define the link station on this host link.

ITSO/3746-9x0/ ESCON Station Configuration - Port 2240

APPN host link: IP host link: HL2240IP Number of host links: 1

Example: Stations

Host or Partitions
VTAM/TPF VTAM/TPF
Unit Add: Unit Add:
01 02 03



Configure an ESCON Station

Network? ☐ APPN (A) ☒ IP (I) ☐ SNA/Subarea (S)
☐ VTAM ☐ TPF Name: S224016
PU type: ☒ 1 ☐ 2.1 ☐ 5 Unit Address (UA): 10 hex
IPL through that station? ☐ Yes ☒ No
On which CCU? ☐ CCU-A ☒ CCU-B
IP address: 192.168.202.8 IP subnet mask: 255.255.255.0
Comments (optional):

ESCON Stations Already Configured

Name	Network	PU	UA	CCU
S224016	IP	1	10	No

Buttons: Add, Modify, Delete, DLC parameters..., APPN parameters..., OK, Cancel, Help

Figure 290. ESCON Station Configuration Panel - Port 2240

The IP Station parameter settings are illustrated in Figure 290. The IP name is set to S224016 with a unit address equal of 10 (hexadecimal). This means in our case the unit address 0x10 refers to 0x90F in the host IOCP definitions. Furthermore, we specified the IP address 192.168.202.8 using subnet mask 255.255.255.0.




Figure 291. Coupler/Port 2240 Configured

After entering the IP parameters and selecting **OK** we returned to the main CCM configuration menu. Note that after configuring the port, port 2240 shows a small check mark.

7.6.1 MVS Definitions

To enable native IP over ESCON transport between TCP/IP for MVS and the 3746 IP router, requires IOCP and TCP/IP for MVS configuration in addition to the 3746-9x0 definitions. The IOCP definitions are very straightforward as they are exactly the same as used for NCP or 3746 NN/DLUR controlled channel connections.

Note: This simplifies the migration from NCP IP to 3746 IP greatly.

Figure 292 depicts the IOCP definitions used during our testing. The 3746-9x0 was attached to the host system via a single ESCON director (ESCD).

```

*-----
*IOCP   3746 NN Directly attached (NO ESCD)
*-----
          CNTLUNIT CUNUMBR=900,UNIT=3745,PATH=(18),
          CUADD=0,LINK=(EA),UNITADD=((01,32))
DEV900   IODEVICE UNIT=3745,ADDRESS=(900,32),CUNUMBR=(900),UNITADD=01
  
```

Figure 292. IOCP Definitions for 3746 IP (with ESCD)

In addition to the IOCP definitions, TCP/IP for MVS requires a limited number of definitions for each ESCON CDLC link in the TCP/IP for MVS PROFILE data set. An example is given in Figure 293 on page 307.

```

DEVICE DEVXA8D CDLC 90F 200 200 2060 2060
LINK LINKXA8D CDLC 0 DEVXA8D

HOME
192.168.202.18 LINKXA8D ; 3746 IP CDLC Escon Channel

START DEVXA8C

```

Figure 293. TCP/IP for MVS CDLC Definition

Make sure that the IO address in the DEVICE statement (90F) matches your IOCP and 3746 ESCON definitions. In our case, IO address 90F (16th address) corresponds with the link station 10 (0x10, 16th station) defined on the 3746.

Figure 294 shows example TCP/IP for MVS static host route and default route definitions.

```

GATEWAY
;
;NETWORK FIRST HOP LINK PACKET SIZE SUBN MASK
192.168.202.18 = LINKXA8D DEFAULTSIZE HOST
DEFAULTNET 192.0.0.0 LINKXA8D DEFAULTSIZE 0

```

Figure 294. Default and Static Route Definitions

7.7 IP Parameter Configuration

After having enabled its ports for IP, further customization of the 3746-9x0 IP functions can be done by selecting the **IP** menu within the Configuration pull-down menu (see Figure 295 on page 308). Options exist to define General parameters, Static routes, Access controls, Filters, and Bootp forwarding.




Figure 295. IP Configuration Options

7.7.1 IP General Parameter




Figure 296. IP General Parameters

The general IP parameters are depicted in Figure 296. All defaults have been accepted with the exception of the Router ID. The Router ID specifies the IP address the 3746 IP routers use when generating IP traffic (for example, PING messages). This address also represents the router itself. For example, this means you can ping the internal router address if there is at least one physical way you can reach the router. If you don't enter one, the code will internally take the lowest IP address defined on an interface. So there is always a router ID, even if you do

not configure it yourself. During our testing we set the Router ID to the 3746 service LAN IP network (192.168.152.99).

Important

During our testing we found that SNMP wouldn't work with this address. An expanded test showed that the Router ID has to be either undefined or the IP address of port 2080 (Service LAN), which is in our test scenario 192.168.152.3.

It is also possible to configure and display the Router ID in a telnet session (see Figure 297).

```
RANGE 2048-2111 *talk 6
Gateway user configuration
Config>protocol ip
Internet protocol user configuration
IP config>list addresses
IP addresses for each interface:
  0 PORT2080 192.168.152.3 255.255.255.0 Local wire broadcast, fi
  1 PORT2144 192.168.150.3 255.255.255.0 Local wire broadcast, fi
                9.24.104.168 255.255.255.0 Local wire broadcast, fi
  2 PORT2176 192.168.157.2 255.255.255.0 Local wire broadcast, fi
  3 PORT2208 10.0.0.3 255.0.0.0 Local wire broadcast, fi
  4 P2240IP 192.168.202.8 255.255.255.0 Local wire broadcast, fi
  5 PP2310IP 0.0.0.5 0.0.0.0 Local wire broadcast, fi
  6 PP2311IP 192.168.155.9 255.255.255.0 Local wire broadcast, fi
  7 X25IP 192.168.159.9 255.255.255.0 Local wire broadcast, fi
Router-ID: 192.168.152.99
Internal IP address: 192.168.152.99
IP config>exit
Config>exit

RANGE 2048-2111 *
```

Figure 297. Display the Internal IP Address and the Router ID

If you define the Router ID in CCM, both the Router ID and the Internal IP Address are set to this value. Only when using the Telnet interface can you specify a separate address for each. Referring to 7.12, “SNMP Parameter Configuration” on page 326, the internal IP address has to be 0.0.0.0 or the address of Port 2080. The Router ID may have any IP address.

7.7.2 PPP NCP Definition

Figure 298 on page 310 shows the dialog where the PPP NCP defaults can be defined. For more information about PPP and NCP see 5.8, “Point-to-Point Protocol (PPP)” on page 193.

ITS0/3746-9x0/PPP NCP Default Parameters

Retry timer: 3000 milliseconds [200-30000]

Config tries: 20 numerical [1-100]

NAK tries: 10 numerical [1-100]

Terminate tries: 10 numerical [1-100]

OK Cancel Help

Figure 298. PPP NCP Definition

7.7.3 IP Static Routes

ITS0/3746-9x0/IP Static Routes

Configure a Static Route

☐ Default route

Destination network: 192.168.156.0

Destination mask: 255.255.255.0

Next hop addresses Cost [numerical 1-16]

	Next hop addresses	Cost
1	192.168.157.9	2
2		1
3		1
4		1

Static Routes Already Configured

Destination network	Destination mask	Next hop address 1	Cost 1
192.168.156.0	255.255.255.0	192.168.157.9	2

Add Modify Delete

OK Cancel Help

Figure 299. IP Static Routes

Static routes can be added, modified or deleted as indicated in Figure 299. A static route for network 192.168.156 with subnet mask 255.255.255.0 using router 192.168.157.9 as its next hop has been added. You can define up to four routes for each destination. For details about which route is actually being used, see 4.4, “Static Routes” on page 145.

7.7.4 IP Access Controls

ITS0/3746-9x0/IP Access Controls

☒ Enable access control

Configure an Access Control Entry—

Access control type: ☐ Permit/Inclusive (I) ☒ Deny/Exclusive (E)

Source: Network IP address Mask address

9.24.104.78 255.255.255.255

Destination: 192.168.152.0 255.255.255.0

Protocol number: From To

0 255 numerical [0-255]

Port number: From To

0 65535 numerical [0-65535]

Access Control Entries Already Configured—

Type	Source IP/Mask addresses	Destination IP/Mask addresses
E	9.24.104.78 255.255.255.255	192.168.152.0 255.255.255.0
I	0.0.0.0 0.0.0.0	0.0.0.0 0.0.0.0

OK Move up Move down Save as defaults Cancel Help

Figure 300. IP Access Controls

Access controls can be defined to ensure that network resources can only be accessed by authorized users and authorized methods of access (see 5.12.2, “3746-9x0 IP Access Controls” on page 230). Figure 300 illustrates the parameters available for customization.

The test environment has an access control list defined forbidding source IP address 9.24.104.78 access to the service LAN subnet (192.168.152). The access control list has the following entries:

- A Denied/Exclusive entry for source IP address 9.24.104.78 to deny access to subnet 192.168.152.
- A Permit/Inclusive entry to allow all remaining traffic to be forwarded.

Note that the last entry in the list cannot be deleted or changed from CCM. As the entries contained in the access control list are interpreted from top to bottom, this means all frames that reach the last entry will be forwarded.

Because we have changed the access to the Service LAN we get a warning (see Figure 301 on page 312).

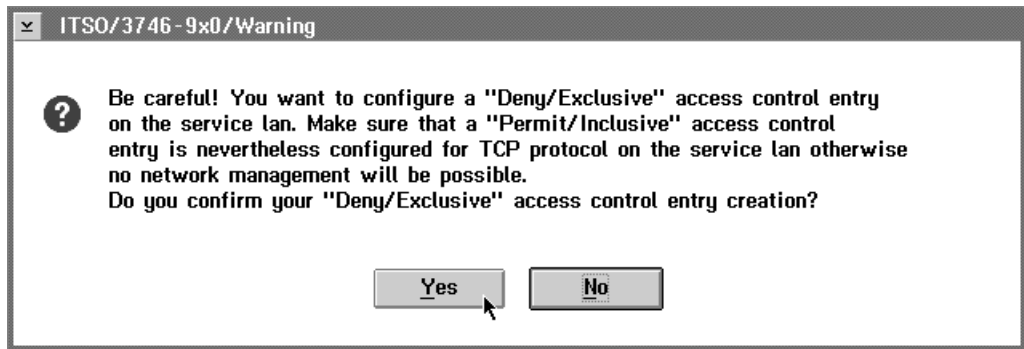


Figure 301. Warning Message

7.7.5 IP Filters

See 5.12.1, “3746-9x0 IP Filters” on page 229 for a description of IP filters. Figure 302 illustrates the filter customization capability to disable routing on the basis of destination IP addresses and subnet masks. In our testing we did not allow IP forwarding to destination subnet 192.168.159.0 via the 3746 IP router.




Figure 302. IP Filters Panel

7.7.6 IP BootP Forwarding




Figure 303. Bootp Forwarding

Figure 303 illustrates Bootp forwarding customization. Bootp is a bootstrap protocol used by a router or a diskless workstation to learn its IP address and obtain a boot file. If Bootp forwarding is to be defined, add the Bootp server address(es) and select the Enable bootp forwarding option.

7.8 OSPF Parameter Configuration

In our test environment we have activated OSPF functions on port 2176. OSPF is an interior routing protocol that maintains link-state information concerning various networks within the OSPF configured autonomous system (AS).




Figure 304. OSPF Configuration Options

CCM customization of OSPF parameters is possible via the OSPF option in the Configuration pull-down menu (see Figure 304).

Options exist to configure General/Multicast routing control, Parameters per IP address, Areas, Imported/default routes, and Virtual links.

ITSO/3746-9x0/OSPF - General/Multicast Parameters

General OSPF Parameters

☒ Enable OSPF

Number of AS external routes: 100 numerical [1-4000]

Number of OSPF routers: 50 numerical [1-4000]

Multicast

☐ Enable intra-area multicasting

☐ Enable inter-area multicasting

Group Addresses

Group address: []

[]

Add

Modify

Delete

OK Cancel Help

Figure 305. OSPF - General/Multicast Parameters

The general OSPF parameters are depicted in Figure 305. We enabled OSPF and used default values for the other parameters.




Figure 306. OSPF/RIP - Parameters per IP Address

OSPF can be configured per IP address as illustrated in Figure 306. Our test environment has token-ring port 2176 with IP address 192.168.157.2 configured for OSPF by selecting **Add OSPF**. Specific OSPF parameters can be entered after selecting **OSPF parameters**.




Figure 307. OSPF - Parameters per IP Address

Figure 307 illustrates the OSPF default parameters defined for port 2176 IP address 192.168.157.2. Once all applicable options are defined, select **OK**.

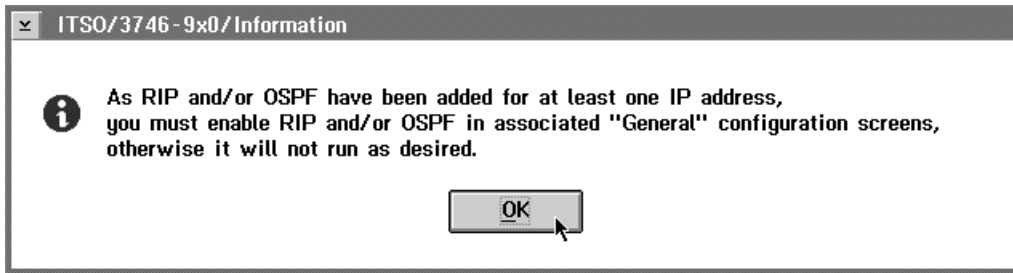


Figure 308. OSPF Enable Warning

Figure 308 occurs upon completion of OSPF parameters per IP address. As mentioned, Enable OSPF must be selected in Figure 305 on page 314 to enable the use of OSPF.




Figure 309. OSPF - Areas Configuration

Figure 309 depicts how OSPF area configuration can be entered. Our test environment contains only a single area, the OSPF backbone area 0.0.0.0. Selecting the Ranges option enables OSPF area border routers to aggregate routes and export a single or a limited number of routes for each area. This facility is very useful in large networks to limit the size of routing tables. For our small test network this function has no value.




Figure 310. OSPF - Area Ranges

Figure 310 illustrates the ability to add, modify or delete OSPF area ranges and specify their associated IP addresses and subnet masks. As mentioned, we did not use the route aggregation function in our test network.




Figure 311. OSPF - Imported/Defaulted Routes

OSPF imported and default routes configuration is possible from the OSPF option in the Configuration pull-down menu (see Figure 304 on page 313). The different route protocols imported are available after selecting the **Enable routes import** option. For our test network we import all known routes.




Figure 312. OSPF - Virtual Links

OSPF virtual links can be specified via the OSPF option in the Configuration pull-down menu (see in Figure 304 on page 313). Virtual links are used to ensure that the OSPF backbone remains connected. A virtual link is a type of PPP link that connects two area border routers that are separated by a non-backbone transit area. Our small, test environment has no OSPF virtual links defined.

7.9 RIP Parameter Configuration

Within our test network we activated RIP functions on port 2240.




Figure 313. RIP Configuration Options

CCM customization of RIP parameters is available from the RIP option in the Configuration pull-down menu (see Figure 313). Options exist to define General RIP parameters and Parameters per IP address.




Figure 314. RIP - General Parameters

The general RIP parameters are depicted in Figure 314. The default parameter values have been taken, except that we enabled RIP.




Figure 315. OSPF/RIP - Parameter per IP Address

RIP can be configured per IP address as illustrated in Figure 315. In our test environment we use RIP on port 2240 with IP address 192.168.202.8. Having selected RIP for this port, additional RIP parameters can be defined by selecting RIP parameters.




Figure 316. RIP - Parameter per IP Address

Figure 316 illustrates the RIP default parameters defined for IP address 192.168.202.8. Select **OK** to accept the parameters.




Figure 317. RIP Enable Warning

Figure 317 occurs upon completion of RIP parameters per IP address. As implied, enabling RIP is essential (see Figure 314 on page 319) to use RIP on the 3746 IP router.

7.10 BGP Parameter Configuration

In our test environment we defined the use of BGP on port 2208. BGP is an exterior routing protocol.




Figure 318. BGP Configuration Options

CCM customization of BGP parameters is available from the BGP option in the Configuration pull-down menu (see Figure 318). As indicated, we then have the ability to define General/excluded AS routing control, Receive policies, Send policies, Originate policies, Neighbors, and Aggregate routes.

ITSO/3746-9x0/BGP - General Parameters/Excluded Autonomous Systems (AS)

General Parameters

☒ Enable BGP ☐ Send subnet routes

AS number: 1 numerical [1-65535]

TCP segment size: 1024 bytes [28-65535]

Excluded Autonomous Systems (AS)

AS no. to exclude: 1 numerical [1-65535]

Buttons: Add, Modify, Delete

Buttons: OK, Cancel, Help

Figure 319. BGP - General Parameters

The general BGP parameters are depicted in Figure 319. The defaults have been taken, except that we enabled BGP.

ITSO/3746-9x0/BGP - Neighbors

Configure a BGP Neighbor

IP address: 10.0.0.8 ☒ Enable neighbor

AS number: 2 numerical [1-65535]

Initialization timer: 12 seconds [0-65535]

Connect retry timer: 120 seconds [0-65535]

Hold timer: 90 seconds [0-65535]

TCP segment size: 1024 bytes [28-65535]

Buttons: Add, Modify

BGP Neighbors Already Configured

IP address	Status	AS number
10.0.0.8	Enable	2

Buttons: Delete

Buttons: OK, Cancel, Help

Figure 320. BGP - Neighbors

Figure 320 illustrates the BGP neighbors defined for our test environment. We included the IP address of the 2210 router (10.0.0.8) that is also attached to the BGP LAN. We use AS number 2 for the 2210. All other parameters have been set to their default values.

ITS0/3746-9x0/BGP - Originate Policies

Configure a BGP Originate Policy

Policy Type: ☐ Inclusive (I) ☒ Exclusive (E)

Address Match (AM): ☐ Exact (E) ☒ Range (R)

Network IP address: 0.0.0.0 Network mask: 0.0.0.0 Tag: 0 num. [0-65535]

BGP Originate Policies Already Configured

No.	Type	AM	Network IP address	Network mask	Tag
1	E	R	0.0.0.0	0.0.0.0	0

Buttons: Add, Modify, Delete, Move up, Move down, OK, Save as defaults, Cancel, Help

Figure 321. BGP - Originate Policies

The BGP originate policies decide which interior protocol routes will be advertised to BGP neighbors. Figure 321 illustrates the BGP originate policy defined on the 3746 IP router. We do not allow BGP to advertise any direct, RIP, or OSPF learned route, but instead, using the route aggregation facilities of BGP, advertise a single aggregated route for all 192.168 addresses (see Figure 322). This approach minimizes the number of routes advertised and limits access to hosts in our test network only.

ITS0/3746-9x0/BGP - Aggregate Routes

Aggregate Routes

Network IP address: 192.168.0.0 Network mask: 255.255.0.0

No.	Network IP address	Network mask
1	192.168.0.0	255.255.0.0

Buttons: Add, Modify, Delete, Move up, Move down, OK, Cancel, Help

Figure 322. BGP - Aggregate Routes

Figure 322 defines the BGP aggregate routes. As illustrated, all routes with network address 192.168 in our test environment will be combined and advertised as a single route to BGP neighbors.

ITSO/3746-9x0/BGP - Receive Policies

Configure a BGP Receive Policy

Policy Type: ☒ Inclusive (I) ☐ Exclusive (E)

Address Match (AM): ☐ Exact (E) ☒ Range (R)

Network IP address: 0.0.0.0 Network mask: 0.0.0.0 IGP metric: 1 num. [1-65535]

Originating AS number: 0 num. [0-65535] Adjacent AS number: 0 num. [0-65535]

BGP Receive Policies Already Configured

No.	Type	AM	Network IP address	Network mask	Originating/Adjacent AS numbers
1	I	R	0.0.0.0	0.0.0.0	0 0

Buttons: Add, Modify, Delete, Move up, Move down, OK, Save as defaults, Cancel, Help

Figure 323. BGP - Receive Policies

The BGP receive policies decide which routes learned from BGP will be injected into the interior gateway protocols. Figure 323 illustrates the BGP receive policy customization options. As can be seen we accept any route.

ITSO/3746-9x0/BGP - Send Policies

Configure a BGP Send Policy

Policy Type: ☒ Inclusive (I) ☐ Exclusive (E)

Address Match (AM): ☐ Exact (E) ☒ Range (R)

Network IP address: 0.0.0.0 Network mask: 0.0.0.0 Tag: 0 numerical [0-65535]

Adjacent AS number: 0 numerical [0-65535]

BGP Send Policies Already Configured

No.	Type	AM	Network IP address	Network mask	Tag	Adjacent AS no.
1	I	R	0.0.0.0	0.0.0.0	0	0

Buttons: Add, Modify, Delete, Move up, Move down, OK, Save as defaults, Cancel, Help

Figure 324. BGP - Send Policies

The BGP send policies decide which routes learned from BGP will be advertised to BGP neighbors. Figure 324 illustrates the BGP send policy customization. As can be seen any route learned will be advertised.

7.11 Proxy-ARP Parameter Configuration

The test environment illustrated in Figure 250 on page 283 has no proxy-Address Resolution Protocol (proxy-ARP) defined, and the CCM configuration menus are shown to depict the configuration options available.




Figure 325. ARP Configuration

CCM customization of ARP parameters is available from the ARP option in the Configuration pull-down menu (see Figure 325).

The image shows a window titled "ARP" with a dropdown arrow on the left. It contains three main sections:

- ARP General Parameters:**
 - Three checkboxes: "Enable ARP net routing", "Enable ARP subnet routing", and "Enable auto refresh".
 - A "Refresh timer" set to "5" minutes, with a range of "[0-1000]".
- Configure an ARP Entry:**
 - A "Port number" dropdown menu set to "2080".
 - Fields for "Remote IP address:" and "Remote MAC address:".
 - Buttons: "Add", "Modify", and "Delete".
- ARP Entries Already Configured:**
 - A table with two columns: "Port" and "Remote IP address/MAC address".
 - An empty table body with a vertical scrollbar on the right.

At the bottom are "OK", "Cancel", and "Help" buttons. A mouse cursor is pointing at the "OK" button.

Figure 326. ARP

Figure 326 illustrates the ARP general parameters that can be set for proxy-ARP.

7.12 SNMP Parameter Configuration

The 3746-9x0 SNMP support is described in 5.11, "3746-9x0 SNMP Functions" on page 225. We enabled SNMP access for the network management station with IP address 9.24.104.188, and also configured SNMP TRAPs to be sent to IP address 9.24.104.188.

Important

In our test scenario it turned out that the Router ID (see Figure 296 on page 308) has to be either undefined or the IP address of Port 2080 (Service LAN) which is in our test scenario 192.168.152.3. In any other case SNMP wouldn't work.




Figure 327. SNMP Configuration

CCM customization of SNMP parameters is available using the SNMP option in the Configuration pull-down menu (see Figure 327).




Figure 328. SNMP Read Access

Figure 328 shows how to configure the 3746 so that station 9.24.104.188, using community name *public* (the community name is case sensitive), has read access to the 3746 IP SNMP variables.

SNMP Configuration

Configure a Set of SNMP Transport Information

Access type: ☐ Read ☒ Trapdest

Community name:

UDP transport:

Network IP address: Network mask:

SNMP Transport Information Sets Already Configured

Access	Community name	Network IP address	Network mask
Read	public	9.24.104.188	255.255.255
Trapdest	public	9.24.104.188	

Buttons: Add, Modify, Delete, OK, Cancel, Help

Figure 329. SNMP Traps

Figure 329 show how to configure the 3746 so that the 3746 IP router sends its SNMP TRAPs, using community name *public*, to station 9.24.104.188.

Appendix A. TELNET Line Command Overview

This section overviews the line commands available from the 3746 IP line command operation and configuration services. For more details on these commands, see the *Multiprotocol Routing Services V2.1 Software Users Guide*, SC30-3681, this book provides information on the commands supported by the IBM 2210 and 2216. The 3746 supports a subset of the commands supported by the IBM 2210 and 2216.

A.1 OPCON Commands

The OPCON commands are depicted in Table 9. The commands become available after the OPCON prompt has appeared.

RANGE 2048-2111*

Table 9. OPCON Commands			
Command	Description	Parameters	Description
RAnge	Redirect OPCON commands	portname port # inf #	Redirect OPCON commands
Memory	Display memory usage		Display memory usage for current processor
Restart	Restart 3746-9x0 router processes		Restart 3746-9x0 IP functions
Status	Display status of 3746 processor IP processes	portname port # inf #	Display IP processes on processor indicated
Talk	Enter GWCON or CONFIG command level	5 6	Process ID of CONFIG(6) or GWCON(5)

Notes:

1. The portname is the alphanumeric name assigned by CCM.
2. The port # is the numeric port number (for example, 2144).
3. The interface (inf) # is an internal numeric network interface identifier.

A.2 GWCON Commands

The GWCON commands are depicted in Table 10 on page 330. The commands become available after the GWCON command level has been entered.

RANGE 2048-2111***talk 5**
RANGE 2048-2111+

Table 10. GWCON Commands

Command	Description	Parameters	Description
Buffer	Display buffer usage	portname port # inf #	Display for current processor. Redirect OPCON commands (implicit) and display for alternate processor.
Clear	Clear statistic	portname port # inf #	Display current processor. Redirect OPCON commands (implicit) and display for alternate processor.
CONfig	Display current protocols and interfaces	- list devices - configuration - patches	- network interfaces - supported protocols - patches
DISable	Take interface offline	portname port # inf #	Interface nr (implicit redirect to processor)
ERROR	Show error counts to either GWCON or CONFIG	portname port # inf #	Display for current processor. Redirect OPCON commands (implicit) and display for alternate processor.
EVENT			
Interface	Display network hardware statistics		
LOG			
MEMory	Display memory, buffer, and packet data		
NETwork	Enters console environment of specific network	portname port # inf #	Interface nr (implicit redirect to processor)
PRotocol	Enters console environment of specific network	IP OSPF BGP ARP	Protocol identifier
QQueue	Display buffer statistic for specific interface	portname port # inf #	Interface nr (implicit redirect to processor)
RAnge	Redirect GWCON commands	portname port # inf #	Redirect to other processor
Statistic	Display statistic for specific interface	portname port # inf #	Interface nr (implicit redirect to processor)
TEST	Enables a specific interface or tests the specific interface	portname port # inf #	Interface nr (implicit redirect to processor)
UPtime	Display statistics for router		
ELS-Display			
DEbug	Enter debug mode for processor		
PHDUMP	Take physical dump of processor		
Exit	Return to OPCON		[CTRL] P can be used as well

Notes:

1. The portname is the alphanumeric name assigned by CCM.
2. The port # is the numeric port number (for example, 2144).
3. The interface (inf) # is an internal numeric network interface identifier.

A.3 Config Commands

The Config commands are depicted in Table 11. The commands become available after the Config command level has been entered.

```
RANGE 2048-2111*talk 6
Config>
```

Table 11. CONFIG Commands			
Command	Description	Parameters	Description
Add		V25.BIS Address	Add network address name for each serial line
Event	Event logging system configuration		
Feature	Enter feature configuration mode	BRS	Configure BRS feature
List	Display system parameters or hardware configuration	Devices, Configuration, Patches	
Network	Enter network interface configuration environment	portname port # inf #	Interface nr (implicit redirect to processor)
PATCH	Modify router global configuration	• RIP-static-suppress • bgp-subnets • PING-size • PING-ttl • ip-default-ttl	• RIP advertising non-static routes • Export subnets in addition to nets • Default size PING (excluding IP header) • TTL time set in PING • TTL set in packets from 3746 IP itself
PRotocol	Enter protocol configuration environment	IP ARP OSPF BGP	Enter protocol configuration mode
Set	Set system wide parameters		
Time	Keep track of system time		Set/display time
Unpatch	Restores patch variables to default values		
Exit	Return to OPCON		[CTRL] P can be used as well

A.4 IP Monitoring and Configuration

The next two sections overview the TELNET line commands available for IP configuration and monitoring.

A.4.1 Configuring IP

Table 12 depicts the commands available for IP configuration. The commands can be entered after the IP Config prompt has appeared.

```
RANGE 2048-2111*talk 6
Config>protocol IP
IP Config>
```

Table 12 (Page 1 of 2). Configuring IP

Command	Description	Parameters	Description
Add	Add IP attributes	• ACCEPT-Rip-route • ACCESS-control • ADDRESS • Bootp-server • Filter • Route	• Accept RIP route • Adds entry to access list • Assign IP address list • Add Bootp server to list • Filter IP network/subnet • Add static route
Change	Modify IP configuration	• Address • Filter • Router	• Modify interface address • Modify subnet mask for (sub)net • Modify static route
Delete	Delete IP configuration	• ACCEPT-Rip-route • ACCESS-control • ADDRESS • Bootp-server • Default-network • Filter • Route	• Delete RIP route • Delete access list record • Remove IP address for interface • Delete server from list • Delete default (subnet) gateway • Delete filter • Delete static route
Disable	Disable IP attributes	• Arp-subnet-routing • ARP-net-routing • Bootp-forwarding • Directed-broadcast • OVERRIDE... • per-packet-multipath • RECEIVING Rip • RECEIVING... • Rip • SENDING... • SOURCE-ROUTING	• Turn off ARP subnet routing • Turn off ARP networking routing • Turn off Bootp forwarding • Turn off directed-broadcast • No overriding default/static route • Turn off multipath routing • Disable accepting specific (sub)nets • Turn off Rip • Turn off RIP sending default, (sub)net, or static routes, poisoned reverse routes • Deactivate support for source-routing
Enable	Enable IP attributes	• Arp-subnet-routing • ARP-net-routing • Bootp-forwarding • Directed-broadcast • OVERRIDE... • per-packet-multipath • RECEIVING Rip • RECEIVING... • Rip • SENDING... • SOURCE-ROUTING	• Turn on proxy-ARP • Turn on proxy-ARP • Turn on Bootp forwarding • Turn on directed-broadcast • No overriding default/static route • Turn on multipath routing • Turn on RIP on interface • Enable RIP accepting specific (sub)net, or static routes, poisoned reverse routes • Activate support for source-routing

Table 12 (Page 2 of 2). Configuring IP			
Command	Description	Parameters	Description
List	Display IP attributes	• ALL • ACcess-controls • ADdress • Bootp • Protocols • Rip-routes-accept • ROutes • Sizes • TAGs	• Display entire IP configuration • Display access control list • Display Bootp servers • Display addresses/broadcast type • Display status IP routing protocols • Display always accepted routes • Display static routes • Size routing table, reassembly buffer size, and route cache • Per-interface TAG associated with received RIP information
Move	Reorder access control list	access-control #1 #2	Move record#1 after #2
Set	Establishes IP configuration modes	• ACcess-control • Broadcast-address • Cache-size • DEFAULT... • IInternal-ip-address • originate-rip-default • REassembly-size • ROuting • ROUTER-ID • Routing table size • TAG • TTL	• Enable/disable access control • Set type of broadcast • Set maximum # in routing cache • Set default (sub)net router • Set internal IP address (Use 0.0.0.0) • RIP to advertise default route • Size of buffers used for reassembly • Size of routing table • Default IP address • Source address for PING, TRACEROUTE, OSPF • Per-interface TAG associated with received RIP information • Time-to-live for 3746 IP datagrams
Exit	Exit IP configuration		[CTRL] P can be used as well

Note: The adding/deleting of static routes and IP filters does not require a router restart, as these commands are executed immediately.

A.4.2 Monitoring IP

Table 13 depicts the commands available for IP monitoring. The commands can be entered after the IP prompt has appeared.

```
RANGE 2048-2111*talk 5
RANGE 2048-2111+protocol IP
IP>
```

Table 13 (Page 1 of 2). Monitoring IP			
Command	Description	Parameters	Description
RANGE	Redirect OPCON commands	portname port # inf #	Redirect OPCON commands
access	List access control mode/records		
cache	Display IP routing cache		
counters	Display statistics related to IP forwarding process		
dump	Display whole IP routing table		
interface	Display interfaces IP addresses and masks		

Table 13 (Page 2 of 2). Monitoring IP			
Command	Description	Parameters	Description
PACKET-FILTER	Display packet filters		
PING	Verify routes to destination address	IP address	• source address optional parameter • stop PING with ENTER
route	Display routes to destination address	IP address	Destination address
Sizes	Display configured sizes of specific IP parameters		
static	Display static routes configured		
Traceroute	Display complete path to destination address	IP address	Stop Traceroute with ENTER
Exit	Exit IP monitoring		[CTRL] P can be used as well

A.5 OSPF Monitoring and Configuration

The next two sections overview the TELNET line commands available for OSPF configuration and monitoring.

A.5.1 Configuring OSPF

Table 14 depicts the commands available for OSPF configuration. The commands can be entered after the OSPF Config prompt has appeared.

```
RANGE 2048-2111*talk 6
Config>protocol OSPF
OSPF Config>
```

Table 14 (Page 1 of 2). Configuring OSPF			
Command	Description	Parameters	Description
add	Add OSPF information	• range • neighbor	• Add address range to OSPF area • Add neighbor to non-broadcast network
delete	Delete OSPF configuration	• range • area • interface • neighbor • non-broadcast • virtual-link	• Delete address range • Delete area • Delete interface • Delete neighbor • Delete non-broadcast network info • Delete virtual-link
disable	Disable entire OSPF, AS boundary, or IP multicast routing	• as boundary routing • multicast forwarding • OSPF routing protocol	• Disable route import • Disable IP multicast • Disable OSPF
enable	Enable entire OSPF, AS boundary, or IP multicast routing	• as boundary routing • multicast forwarding • OSPF routing protocol	• Enable route import • Enable IP multicast • Enable OSPF

Table 14 (Page 2 of 2). Configuring OSPF			
Command	Description	Parameters	Description
join	Join multicast group	group address	Configure 3746-9x0 as member of multicast group
leave	Leave multicast groups	group address	Remove membership in multicast group
list	Display OSPF configuration	- all - areas - interfaces - neighbors - non-broadcast - virtual-links	- List all OSPF configuration - List OSPF area - List OSPF interfaces - List neighbors - non-broadcast network info - List virtual links
set	Establishes OSPF configuration modes	- area - comparison - interface - non-broadcast - virtual-links	- Set parameters for OSPF area - Hierarchy of BGP/RIP/static routes - Set OSPF parameters for interface - Set OSPF neighbors - Define virtual link
Exit	Exit OSPF configuration		[CTRL] P can be used as well

Note: The adding/deleting of OSPF neighbors does not require a router restart, as these commands are executed immediately.

A.5.2 Monitoring OSPF

Table 15 depicts the commands available for OSPF monitoring. The commands can be entered after the OSPF prompt has appeared.

```
RANGE 2048-2111+talk 5
RANGE 2048-2111+protocol OSPF
OSPF>
```

Table 15 (Page 1 of 2). Monitoring OSPF			
Command	Description	Parameters	Description
RANGE	Redirect OPCON commands	portname port # inf #	Redirect OPCON commands
ADVer-tisement	Display a link-state advertisement belonging to OSPF database	ls-type	Display LSAs, per link stat ID, per advertising router, per area
area	Display OSPF areas information		
AS-external	List AS external advertisements		
database	Display advertisement belonging to an area's link-state database	area-ID	List for particular area ID
dump	Display whole routing table		
interface	Display OSPF interface information	IP-address	List OSPF info for 3746-9x0 interface
join	Add router to multicast group	IP-address	Join multicast group
leave	Leave multicast group	IP-address	Leave multicast group

Table 15 (Page 2 of 2). Monitoring OSPF

Command	Description	Parameters	Description
mcache	Display active multicast forwarding cache entries	source dest	Display overview, or detail for source and destination IP addresses
mgroups	Display group membership of router interfaces		
mstat	Display OSPF multicast routing statistics		
neighbor	Display OSPF neighbor statistics	IP-address	Display information for neighbor
PING	Test reachability destination	IP-address	PING to IP-address (ENTER to stop)
routers	Display OSPF area-border and AS-boundary routers		
size	Display number of LSAs in link-state database		
statistics	Display OSPF statistics (memory/network usage)		
Traceroute	Display path to destination	IP-address	Traceroute to IP-address (ENTER to stop)
weight	Change cost of OSPF interface	IP-address cost	Change OSPF weight of interface "IP-address" into cost
Exit	Exit from OSPF monitoring		[CTRL] P can be used as well

A.6 BGP Monitoring and Configuration

The next two sections overview the TELNET line commands available for BGP configuration and monitoring.

A.6.1 Configuring BGP

Table 16 on page 337 depicts the commands available for BGP configuration. The commands can be entered after the BGP Config prompt has appeared.

```
RANGE 2048-2111*talk 6
Config>protocol BGP
BGP Config>
```

Table 16. Configuring BGP			
Command	Description	Parameters	Description
Add	Add BGP information	• aggregate • neighbor • no-receive asnum • originate-policy • receive-policy • send-policy	• Aggregate block of addresses for advertisements • Define BGP neighbor • Excludes updates from AS • Advertising policy • Importing policy • Forwarding policy
change	Change BGP configuration	• aggregate • neighbor • originate-policy • receive-policy • send-policy	• Change current aggregate • Change BGP neighbor info • Advertising policy • Importing policy • Forwarding policy
delete	Delete previously added configuration	• aggregate • neighbor • no-receive asnum • originate-policy • receive-policy • send-policy	
disable	Disable BGP neighbor or speaker	• BGP speaker • neighbor	• Disable BPG speaker • Disable BPG neighbor
enable	Enable BGP speaker or neighbor	• BGP speaker • neighbor	• Enable BGP • Enable BGP
list	List BPG configuration configured	• aggregate • all • BGP • neighbor • no-receive • originate-policy • receive-policy • send-policy	• List all BGP parameters • List network aggregates • List BGP neighbors • List discarded routes • List originate policy • List receive policy • List send policy
Move			
Exit	Exit BGP configuration		[CTRL] P can be used as well

A.6.2 Monitoring BGP

Table 17 depicts the commands available for BGP monitoring. The commands can be entered after the BGP prompt has appeared.

```
RANGE 2048-2111*talk 5
RANGE 2048-2111+protocol BGP
BGP>
```

Table 17 (Page 1 of 2). Monitoring BGP			
Command	Description	Parameters	Description
RANGE	Redirect OPCON commands	portname port # inf #	Redirect OPCON commands
destination	Dump BGP routing table, or display information on routes advertised to, or received from, specific BGP neighbors	• net address/mask • advertised-to network • Received-from network	• Display routing info • List routes advertised • List routes received
Dump	Display whole routing table		

Table 17 (Page 2 of 2). Monitoring BGP

Command	Description	Parameters	Description
Neighbors	Display info on all active BGP neighbors	neighbors address	Display neighbor info for particular neighbor
Paths	Display paths stored in path description database		
PING	Test reachability destination	IP-address	PING to IP-address (stop PING with ENTER)
Sizes	Display number of entries stored in various data bases		
Traceroute	Display path to destination	IP-address	Traceroute to IP-address (stop Traceroute with ENTER)
Exit	Exit from BGP monitoring		[CTRL] P can be used as well

A.7 ARP Monitoring and Configuration

The next two sections overview the TELNET line commands available for ARP configuration and monitoring.

A.7.1 Configuring ARP

Table 18 depicts the commands available for ARP configuration. The commands can be entered after the ARP Config prompt has appeared.

```
RANGE 2048-2111*talk 6
Config>protocol ARP
ARP Config>
```

Table 18 (Page 1 of 2). Configuring ARP

Command	Description	Parameters	Description
Add Entry	Add ARP entry	ifc# IP IP-address MAC	Add ARP entry for IP address on interface (ifc) #
Change Entry	Modify ARP entry	ifc# IP IP-address MAC	Change ARP entry for IP address on interface (ifc) #
Delete Entry	Delete ARP entry router processes	ifc# IP IP-address MAC	Delete ARP entry
DISable Auto-refresh	Disable ARP auto-refresh		
Enable Auto-refresh	Enable ARP auto-refresh		
List	Display ARP configuration	- All - Config - Entry	- Display all ARP information - List configuration for different timers - List ARP entries

Table 18 (Page 2 of 2). Configuring ARP			
Command	Description	Parameters	Description
Set	Set ARP usage and refreshes timeout values	Refresh-timer	Set refresh-timer, Enter value in minutes, '0' indicates NO refresh
Exit	Exit ARP configuration		[CTRL] P can be used as well

A.7.2 Monitoring ARP

Table 18 on page 338 depicts the commands available for ARP monitoring. The commands can be entered after the ARP prompt has appeared.

```
RANGE 2048-2111*talk 5
RANGE 2048-2111+protocol ARP
ARP>
```

Table 19. Monitoring ARP			
Command	Description	Parameters	Description
RANge	Redirect OPCON commands	portname port # inf #	Redirect OPCON commands
Clear	Flush ARP cache	ifc#	Flush ARP cache for interface (ifc) #
Dump	Display IP routing cache	ifc# IP	Display IP ARP cache for interface (ifc) #
hardware	Display interfaces registered with ARP		
PING	Test reachability destination	IP-address	PING to IP-address (stop PING with ENTER)
Protocol	Display protocols registered per interface		
Statistics	Display statistical information		
Exit	Return from monitoring ARP		[CTRL] P can be used as well

A.8 BRS Monitoring and Configuration

The next two sections overview the TELNET line commands available for BRS configuration and monitoring.

A.8.1 Configuring BRS

Table 20 on page 340 depicts the commands available for BRS configuration. The commands can be entered after the BRS Config prompt has appeared.

```
RANGE 2048-2111*talk 6
Config>feature BRS
BRS Config>
```

Table 20. Configuring BRS

Command	Description	Parameters	Description
interface	Select interface to which BRS commands apply	inf#	Interface #
list	List currently defined bandwidth classes	inf#	Interface #
Exit	Exit BRS configuration		[CTRL] P can be used as well

A.8.2 Monitoring BRS

Table 21 depicts the commands available for BRS monitoring. The commands can be entered after the Reservation console prompt has appeared.

```
RANGE 2048-2111*talk 5
RANGE 2048-2111+feature BRS
Reservation console>
```

Table 21. Monitoring BRS

Command	Description	Parameters	Description
Clear	Clear current counters		
counters	Display counters		
interface	Select interface for BRS console commands	ifc# (dlci#)	Interface nr plus (optional) DLCI nr if frame relay circuit. Note: This command must be entered first.
last	Display last saved counters		
Exit	Exit BRS monitoring		[CTRL] P can be used as well

A.9 ESCON Port Monitoring and Configuration

The next two sections overview the TELNET line commands available for ESCON port configuration and monitoring.

A.9.1 Configuring ESCON

Table 20 on page 340 depicts the commands available for ESCON port configuration. The commands can be entered after the ESCON Config prompt has appeared.

```
RANGE 2048-2111*talk 6
Config>network xxxx      (xxxx = ESCON port)
ESCON Config>
```

Table 22. Configuring ESCON			
Command	Description	Parameters	Description
List	List configuration for interface		
Packet-size	Change packet size for ESCON interface	nnn	
Exit	Exit ESCON configuration		[CTRL] P can be used as well

A.9.2 Monitoring ESCON

Table 23 depicts the commands available for ESCON port monitoring. The commands can be entered after the ESCON prompt has appeared.

```
RANGE 2048-2111*talk 5
RANGE 2048-2111+network xxxx      (xxxx = ESCON port)
ESCON>
```

Table 23. Monitoring ESCON			
Command	Description	Parameters	Description
Dump	Display routing table		
Exit	Exit ESCON monitoring		[CTRL] P can be used as well

A.10 Token-Ring Port Monitoring and Configuration

The next two sections overview the TELNET line commands available for token-ring port configuration and monitoring.

A.10.1 Configuring Token-Ring

Table 24 on page 342 depicts the commands available for token-ring port configuration. The commands can be entered after the TKR Config prompt has appeared.

```

RANGE 2048-2111*talk 6
Config>network xxxx      (xxxx = token-ring port)
TKR Config>

```

Table 24. Configuring IEEE 802.5 Token-Ring

Command	Description	Parameters	Description
SOURCE-ROUTING	Enable or disable source routing	ENABLE DISABLE	
List	List configuration for interface		
Packet-size	Change packet size for token-ring interfaces	nnn	4 Mbps - use 1500,2052, 4472 16 Mbps - use 1470, 2052, 4399, 8130, 11407, or 17749
SET			
Exit	Exit TKR configuration		[CTRL] P can be used as well

A.10.2 Monitoring Token-Ring

Table 25 depicts the commands available for token-ring port monitoring. The commands can be entered after the TKR prompt has appeared.

```

RANGE 2048-2111*talk 5
RANGE 2048-2111*network xxxx      (xxxx = token-ring port)
TKR>

```

Table 25. Monitoring IEEE 802.5 Token-Ring

Command	Description	Parameters	Description
Dump	Display entries in RIF cache		Requires source-routing enabled
RAngE	Redirect GWCON commands	portname port # inf #	Redirect GWCON commands
Exit	Exit ESCON monitoring		[CTRL] P can be used as well

A.11 Frame Relay Port Monitoring and Configuration

The next two sections overview the TELNET line commands available for frame relay port configuration and monitoring.

A.11.1 Configuring Frame Relay

Table 26 on page 343 depicts the commands available for frame relay port configuration. The commands can be entered after the FR Config prompt has appeared.

```

RANGE 2048-2111*talk 6
Config>network xxxx      (xxxx = frame relay port)
FR Config>

```

Table 26. Configuring Frame Relay

Command	Description	Parameters	Description
Add	Add a protocol	Protocol-address	Type of protocol
Change	Change PVC	Permanent-VC	change PVC #, committed information rate (CIR), and circuit name
Disable	Disable FR features	• Cir-monitor • Multicast-emulation • Orphan-circuits • Protocol-broadcast	
Enable	Enable FR features	• Cir-monitor • Multicast-emulation • Orphan-circuits • Protocol-broadcast	
List	List configuration for interface	• HDLC • PERmanent-VC • PRotocol-address	• List HDLC configuration • Display configured PVCs • Display static IP addresses
Remove	Add a protocol	Protocol-address	Type of protocol
Exit	Exit FR configuration		[CTRL] P can be used as well

A.11.2 Monitoring Frame Relay

Table 27 depicts the commands available for frame relay port monitoring. The commands can be entered after the FR prompt has appeared.

```

RANGE 2048-2111*talk 5
RANGE 2048-2111+network xxxx      (xxxx = frame relay port)
FR>

```

Table 27. Monitoring Frame Relay

Command	Description	Parameters	Description
Clear	Clear all statistical information		
List	List statistics	• All • Circuit # • Permanent-VC	• Display all information • List info/counters for PVC ## • List info for all PVCs
Chkpt	Lists checkpoints		
Exit	Exit FR monitoring		[CTRL] P can be used as well

A.12 PPP Port Monitoring and Configuration

The next two sections overview the TELNET line commands available for PPP port configuration and monitoring.

A.12.1 Configuring PPP

Table 28 depicts the commands available for PPP port configuration. The commands can be entered after the PPP Config prompt has appeared.

```
RANGE 2048-2111*talk 6
Config>network xxxx      (xxxx = PPP port)
PPP Config>
```

Table 28. Configuring PPP

Command	Description	Parameters	Description
List	List configuration for interface	- All - HDLC - IPCP - LCP - Parameter	- List all - List HDLC parameters - List IPCP options - List parameters for LCP options - List NCP options
Set	Set HDLC, LCP, IPCP, and NCP options and parameters	- IPCP - Parameters - LCP	- Set IPCP parameters (header compression) - Set Link Control Protocol options/parameters (MTU, magic nr yes/no for loopback detection) (config tries #, nak tries #, terminate tries #, retry time) - Set all Network Control Protocol parameters (config tries #, nak tries #, terminate tries #, retry time)
Exit	Exit PPP configuration		[CTRL] P can be used as well

A.12.2 Monitoring PPP

Table 29 depicts the commands available for PPP port monitoring. The commands can be entered after the PPP prompt has appeared.

```
RANGE 2048-2111*talk 5
RANGE 2048-2111+network xxxx      (xxxx = PPP port)
PPP>
```

Table 29 (Page 1 of 2). Monitoring PPP

Command	Description	Parameters	Description
RAnge	Redirect GWCON commands	portname port # inf #	Redirect GWCON commands
Clear	Clear all statistical information		

Table 29 (Page 2 of 2). Monitoring PPP

Command	Description	Parameters	Description
List	Display information related to PPP interface, options, and parameters	• All • CCP • Control LCP • Control IPCP • Control CCP • Errors • IP • IPCP • LCP	• Display all • Compression statistics • List LCP (link control protocol) • List IPCP (IP control protocol) • List compression control • List error information • List IP info over PPP • List IPCP info over PPP • List LCP statistics
LLC			
Exit	Exit PPP monitoring		[CTRL] P can be used as well

Appendix B. 3746 IP Functions - Overview

The following sections overview the functions supported on the 3746 IP router.

B.1 TCP/IP Functions

- RFC 793 - TCP
- RFC 768 - UDP
- RFC 791 - IP
 - IP options supported:
 - Strict source routing
 - Loose source routing
 - Record route
 - Time stamp
 - RFC 1009 - Internet gateways
 - RFC 1112 - IP multicast
 - RFC 950 - Subnetting
 - IP variable-length subnet
 - IP packet filter
 - Source IP address
 - Destination IP address
 - IP protocol number
 - TCP/UDP port number
 - Multiple IP addresses per interface
 - RFC 1191 - Path MTU discovery
 - RFC 919/RFC 922 - Broadcast including:
 - Configurable broadcast address
 - Directed broadcast
 - UDP broadcast
 - RFC 951/RFC 1542 - Bootp relay agent
- RFC 826/RFC 1942 - ARP
 - RFC 1027 - Proxy-ARP (subnet routing)
 - RFC 925 - ARP subnet routing
 - ARP cache manipulation
 - Configurable aging timer
 - Automatic refresh option
- ICMP
 - Echo request/reply (PING)
 - Destination unreachable
 - Source redirect
 - Time exceeded
 - Address mask request/response

B.2 Dynamic Routing Protocols

- RFC 1058 - RIP Version 1
 - RIP policies:
 - Route acceptance
 - Route advertisement
 - Interaction with other routing protocols
 - Default origination
 - RIP net/subnet/host routes
- RFC 1583/1584 - OSPF Version 2 + Multicast extensions
 - Splitting the AS into areas
 - Stub area support
 - IP subnetting support
 - Point-to-point network support
 - Point-to-multipoint network support
 - Broadcast network support
 - Non-broadcast network support
 - Virtual-links
 - Authentication of neighbors
 - Unnumbered serial line support
 - Equal-cost multipath (load balancing)
 - MOSPF-IP multicast routing support
 - IP supernetting and subnet 0 support
- RFC 1654 - BGP Version 4

B.3 Interfaces Supported

- ESCON, using native IP over CDLC
- Token-ring
- Ethernet
- RFC 1331/1332 - PPP
 - RFC 1144 - TCP/IP header compression
 - Maximum RU and magic number negotiated
 - Bandwidth reservation (BRS)
 - Leased line: V.24, V.35, X.21
- RFC 1490 - frame relay

B.4 SNMP Support

- RFC 1155 - Structure of management information
- RFC 1147, RFC 1441-1450, RFC 1452 - SNMP
- RFC 1592 - SNMP distributed programming interface
- MIBs supported:
 - RFC 1213 - MIB II
 - RFC 1253 - OSPF V2 MIB
 - RFC 1231 - token-ring MIB
 - RFC 1315 - frame relay MIB

- RFC 1471 - PPP IP MIB
- ESCON MIB (IBM-proprietary)

Appendix C. Special Notices

This publication is intended to help system engineers, system planners, system programmers, and network administrators who need to understand stand-alone IP functions offered by the 3746 Model 900 and the 3746 Nways Controller. The information in this publication is not intended as the specification of any programming interfaces that are provided by the 3746 Model 900 and the 3746 Nways Controller. See the PUBLICATIONS section of the IBM Programming Announcement for the 3746 Model 900 and the 3746 Nways Controller for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating

environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ACF/VTAM	Advanced Peer-to-Peer Networking
AIX	AIX/6000
AnyNet	APPN
AS/400	DB2
ESCON	IBM
IBM Global Network	MVS
NetView	Nways
OS/390	OS/2
RISC System/6000	S/390
SP	System/390
VTAM	IBM®

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names may be trademarks or service marks of others.

Appendix D. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 355.

- *IBM 3746 Nways Controller Model 950 and IBM Model 900 - APPN Implementation Guide*, SG24-2536
- *IBM Frame Relay Guide*, GG24-4463
- *TCP/IP Tutorial and Technical Overview*, GG24-3376
- *The Basics of IP Network Design*, SG24-2580

D.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

D.3 Other Publications

These publications are also relevant as further information sources.

- *IBM 3746 Nways Multiprotocol Controller Models 900 and 950 Overview*, GA33-0180
- *3745 Models A and 3746 Model 900 Migration and Planning Guide*, GA33-0183
- *IBM 3746 Nways Multiprotocol Controller Models 900 and 950 Migration and Planning Guide*, GA33-0349
- *IBM 3746 Nways Multiprotocol Controller Model 950 Safety Information*, GA33-0400
- *Input/Output Equipment Installation Manual Physical Planning*, GC22-7064
- *3745 All Models Advanced Operation Guide*, SA33-0097
- *3745 Connection and Integration Guide*, SA33-0129
- *IBM 3746 Expansion Unit Model 900 Console Setup Guide*, SA33-0158
- *3745 Master Index*, SA33-0172

- *IBM 3745 and IBM 3746 Alert Reference Guide*, SA33-0175
- *IBM 3746 Expansion Unit Model 900 Basic Operations Guide*, SA33-0177
- *IBM 3746 Nways Multiprotocol Controller Model 950 User's Guide*, SA33-0356
- *3745 Maintenance Information Procedures*, SY33-2054
- *3745 Service Functions*, SY33-2055
- *3745 Maintenance Information Reference*, SY33-2056
- *3745 Installation Guide*, SY33-2057
- *3745 Service Master Index*, SY33-2080
- *3746 Model 900 Installation Guide*, SY33-2088
- *IBM 3746 Nways Multiprotocol Controller Model 950 Service Guide*, SY33-2108

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at <http://www.redbooks.ibm.com/>.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

To get LIST3820s of redbooks, type one of the following commands:

```
TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
```

To get BookManager BOOKs of redbooks, type the following command:

```
TOOLCAT REDBOOKS
```

To get lists of redbooks, type the following command:

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
```

To register for information on workshops, residencies, and redbooks, type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1998
```

For a list of product area specialists in the ITSO: type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
```

- **Redbooks Web Site on the World Wide Web**
<http://w3.itso.ibm.com/redbooks/>
- **IBM Direct Publications Catalog on the World Wide Web**
<http://www.elink.ibm.link.ibm.com/pb1/pb1>

IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibm.link.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** — send orders to:

	IBMMAIL	Internet
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications Publications Customer Support P.O. Box 29570 Raleigh, NC 27626-0570 USA	IBM Publications 144-4th Avenue, S.W. Calgary, Alberta T2P 3N5 Canada	IBM Direct Services Sortemosevej 21 DK-3450 Allerød Denmark
--	--	--

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States)** or **(+1)001-408-256-5422 (Outside USA)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com

- **On the World Wide Web**

Redbooks Web Site	http://www.redbooks.ibm.com/
IBM Direct Publications Catalog	http://www.elink.ibm.link.ibm.com/pbl/pbl

- **Internet Listserver**

With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibm.link.ibm.com with the keyword `subscribe` in the body of the note (leave the subject line blank).

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity
-------	--------------	----------

First name

Last name

Company

Address

City

Postal code

Country

Telephone number

Telefax number

VAT number

☐ Invoice to customer number

☐ Credit card number

Credit card expiration date

Card issued to

Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

List of Abbreviations

AB	Area border router	CCM	Controller configuration and management program
ARP	Address Resolution Protocol	CCU	Central Control Unit
ACF	Advanced Communications Function	CDF	Configuration Data File
ALS	Adjacent Link Station	CDF-E	Configuration Data File - Extended
ANSI	American National Standards Institute	CDLC	Channel data link control
APA	All Points Addressable	CIR	Committed Information Rate
APPN	Advanced Peer to Peer Networking	CMIP	Common Management Information Protocol
ARB	Adaptive Rate Based	CNN	Composite Network Node
ARC	Active Remote Coupler	CLA	Communication Line Adapter
AS	Autonomous system	CLP	Communication Line Processor
ASB	Autonomous system boundary router	CP	Control Point
ASE	Autonomous system external	CPLR	Coupler
ASM	Address Space Management	CPP	Control Point Processor
ASN.1	Abstract Syntax Notation 1	CRC	Cyclic Redundancy Checking
ATM	Asynchronous Transfer Mode	CSMA/CD	Carrier Sense Multiple Access with Collision Detection
BAS	Basic Board	CR	Communication Rate
BGP	Border gateway protocol	CS	Connectivity Switch
BAT	Basic Assurance Test	CSP	Communication Scanner Processor
BCCA	Buffer Chaining Channel Adapter	CSPRT	Connectivity Subsystem Line Trace Report
BECN	Backward Explicit Congestion Notification	CSS	Communication Subsystem
BER	Box Event Record	CV	Control Vector
BF	Boundary Function	DCAF	Distributed Console Access Facility
CA	Channel Adapter	DCE	Data Circuit Terminating Equipment
CAPRT	Communication Subsystem Adapter Trace Report	DE	Discard Eligibility indicator
CBSP	Controller Bus and Service Processor	DLC	Data Link Control
CBTRP	Controller Bus and Token-Ring Processor	DLCI	Data Link Connection Identifier
CBTRA	Controller Bus and Token-Ring Adapter	DLUR	Dependent LU Requester
CBTRM	Cable Terminator (IOC and DMA buses)	DLUS	Dependent LU Server
CIR	Committed Information Rate	DMA	Direct Memory Access
CCITT	Comite Consultatif International Telegraphique Telephonique	DR	Dynamic Reconfiguration
		DS	Directory Services
		DSAP	Destination Service Access Point

DTE	Data Terminal Equipment	INN	Intermediate Network Node
EAC	Ethernet Adapter Card	IOC	Input/Output Control
EC	Engineering Change	IP	Internet Protocol
EGP	Exterior gateway protocol	IPL	Initial Program Load
ELA	Ethernet LAN Adapter	ISDN	Integrated-Services Digital Network
EMIF	ESCON Multiple Image Facility	ITSC	International Technical Support Center
EN	End Node	ITSO	International Technical Support Organization
EP	Emulation Program	ITU	International Telecommunication Union
EPO	Emergency Power Off	ITU-T	International Telecommunication Union - Telecommunication
ESA	Enterprise System Architecture	KBps	Kilobytes per Second
ESCC	ESCON Coupler	Kbps	Kilobits per Second
ESCD	ESCON Director	LAN	Local Area Network
ESCON	Enterprise Systems Connection	LCB	Line Connection Box
ESCP	ESCON Processor	LCBB	Line Connection Box Base
ESS	Ethernet SubSystem	LCBE	Line Connection Box Expansion
FDX	Full-Duplex	LCD	Link Configuration Data
FECN	Forward Explicit Congestion Notification	LDPRT	Line Trace Detail
FPS	Fast Packet Switching	LEN	Low Entry Networking
FR	Frame Relay	LFSID	Local-Form Session-Identifiers
FRFH	Frame Relay Frame Handler	LIC	Line Interface Couplers
FRSE	Frame Relay Switching Equipment	LIV	Link Integrity Verification
FRTE	Frame Relay Terminal Equipment	LLC	Logical Link Connection
FTP	File transfer protocol	LMI	Local Management Interface
GTF	Generalized Trace Facility	LSA	Link-state advertisement
HDX	Half-Duplex	LPAR	Logical Partition
HPR	High Performance Routing	LSP	Link-state packet
HPTSS	High Performance Transmission Subsystem	LSPRT	Line Trace Summary
IA	inter-area router	LU	Logical Unit
IBM	International Business Machines Corporation	MAC	Medium Access Control
ICMP	Internet Control Message Protocol	MCF	Microcode Fix
IEEE	Institute of Electrical and Electronics Engineers	MCL	Microcode Level
IGP	Interior gateway protocol	MIB	Maintenance information base
IGMP	Internet group management protocol	MLTG	Multi-Link Transmission Group
IML	Initial Microcode Load	MMMLTG	Mixed Media Multi-Link Transmission Group
InARP	Inverse address resolution protocol	MOSPF	Multicast extensions to OSPF

MOSS	Maintenance and Operation Subsystem	RSCV	Route Selection Control Vector
MOSS-E	Maintenance and Operation Subsystem-Extended	RSF	Remote Support Function
MTU	Maximum Transmission Unit	RTP	Rapid Transport Protocol
NBMA	Non-broadcast multi-access	RU	Request Unit
NCP	Network Control Program	SA	Subarea
NDF	Network Definition Facility	SAP	Service Access Point
NDM	NetView Distribution Manager	SC	Session Connector
NEO	Network Extension Option	SCTL	Storage Control
NN	Network Node	SDLC	Synchronous Data Link Control
NNP	Network Node Processor	SINAXE	Secondary Initiated Non-Activation Xid Exchange
NNNA	Non-Native Network Attachment	SIT	Scanner Interface Trace
NMVT	Network Management Vector Transport	SLIP	Serial line interface protocol
NOF	Node Operator Facility	SNA	Systems Network Architecture
NPA	Network Performance Analyzer	SNMP	Simple network management protocol
NPM	NetView Performance Monitor	SP	Service Processor
NPPRT	Line Trace Report	SPS	Service and Power Support
NPSI	NCP Packet Switching Interface	SRB	Source route bridging
NTRI	NCP Token-Ring Interconnect	SRC	System Reference Code
ODLC	Outboard Data Link Control	SS	Session Services
OSPF	Open shortest path first	SSAP	Source Service Access Point
PC	Path Control	SSCP	System Services Control Point
PDU	Protocol Data Unit	SSP	System Support Program
PEP	Partitioned Emulation Program	TAP	Trace Analysis Program
PING	Packet internet groper	TG	Transmission Group
PIU	Path Information Unit	TGN	Transmission Group Number
PPP	Point-to-point protocol	TIC	Token-Ring Interface Coupler
PU	Physical Unit	TOS	Type of service
PVC	Permanent Virtual Circuit	TRA	Token-Ring Adapter
REQMS	Record Maintenance Statistics	TRP	Token-Ring Processor
RARP	Reverse address resolution protocol	TRS	Topology and Route Selection
REXX	Restructured Extended Executor Language	TRSS	Token-Ring Subsystem
RFC	Request for comment	TSS	Transmission Subsystem
RNR	Receiver Not Ready	TTL	Time to live
RIF	Routing information field	UDP	User Datagram Protocol
RIP	Routing information protocol	UEPO	Unit Emergency Power Off
RPS	Route Partner Set	VM	Virtual Machine
		VR	Virtual Route
		VTAM	Virtual Telecommunications Access Method
		WAN	Wide Area Network

<i>X.25</i>	ISO Standard for Interface to Packet Switched Communications Services	<i>XID3 (NP)</i>	Exchange Identification Type 3 (Negotiation Proceeding)
<i>XID</i>	Exchange Identification	<i>XID3 (PN)</i>	Exchange Identification Type 3 (Pre-Negotiation)
<i>XID3 (NA)</i>	Exchange Identification Type 3 (Non-Activation)	<i>XRF</i>	Extended Recovery Facility

Index

Numerics

- 3172 29
 - 3745 162
 - 3745/3746 Overview, GA33-0180-8 30
 - 3745/3746 Planning Guide, GA33-0457 30
 - 3746 APPN CP 186
 - 3746 DLCI sharing 37
 - 3746 hardware addresses 160
 - 3746 IP 2, 185
 - ARP 84
 - BGP 129, 140
 - Bootp 89
 - broadcasting 78
 - configuration 237
 - DLC 55
 - dynamic routing protocols 348
 - ESCON 160
 - ESCON sharing 162
 - Ethernet 175
 - fragmentation 81
 - frame relay 176
 - function overview 347
 - ICMP 65
 - interfaces 160
 - interfaces supported 348
 - IP 54
 - IP address 79
 - IP routing 131
 - operation 237
 - OSPF 118, 134
 - PPP 193
 - proxy-arp 85
 - restart 280
 - RIP 101, 131
 - SNMP 225
 - SNMP support 348
 - TCP 52
 - TCP/IP functions 347
 - token-ring 172
 - token-ring sharing 173
 - UDP 65
 - 3746 IP Management 214
 - 3746 Licensed Internal Code 28
 - 3746 LMI Support 35
 - 3746 Model 900
 - Expansion Unit 2
 - hardware 5
 - Internal Structure 6
 - IP 2
 - NN 2
 - 3746 NN 185
 - 3746 Nways Controller 2
 - connectivity 4
 - hardware 7
 - 3746-900 2
 - 3746-950 2
 - IP 2
 - NN 2
 - 3746-9x0 2
 - connectivity 3
 - 3746-9x0 controller 151
 - 9585 29
- ## A
- abbreviations 359
 - abstract syntax notation 46
 - access control 86, 272, 311
 - access rate 178
 - acronyms 359
 - activate a configuration 250
 - adapter
 - connectivity 19
 - adapters 6
 - add calling DTE Address 213
 - address resolution 81
 - adjacent router 104
 - ANR label 185
 - APING 39
 - application layer 47
 - application programming 56
 - APPN 5, 162, 185, 186, 189
 - APPN network node 7
 - area 102, 316
 - area border router 104
 - area range 317
 - ARP 81
 - ARP table 267
 - AS boundary router 104
 - ASN.1 46
 - ATM adapters 23
 - ATM addressing 24
 - authentication servers 25
 - Autonomous Systems (ASs) Interior gateway protocols 91
- ## B
- backbone 103
 - backup designated router 105
 - backup designated router on the basis of the 109

- backup frame relay 25
- Backward Explicit Congestion Notification (BECN) 37
- Bandwidth Reservation System (BRS) 22, 25
- basic mode 165, 169
- BECN 25, 37, 180
- becoming familiar with the main window 244
- BGP 200
 - 3746 IP 129
 - configuration 141, 321
 - connection 123
 - neighbors 322
 - originate policy 323
 - policies 121
 - protocol description 122
 - receive policy 324
 - route aggregation 323
 - send policy 325
- bibliography 353
- BNN 185, 186, 189
- BNN FRTE 35
- Bootp 88, 313
- Border Gateway Protocol (BGP) 120, 140
- boundary access node (BAN) 37
- boundary function 35
- boundary network node (BNN) 36
- Branch Extender 23, 24
- bridged format 37
- Bridged frame format 185
- broadcasting 77
- BRS 178
- BSD UNIX 99
- burst sizes 178, 179
 - committed 178
 - excess 178
 - Measurement Interval 179

C

- call request 202
- calling DTE address 202, 214
- CASDL 167
- CBC 6
- CBSA 152
- CBSP 152
- CBSP2 6
- CBSP3 6
- CBTRP 6
- CCM 186, 216, 241
 - installing 239
 - IP configuration 283
 - levels 239
 - password 241
 - product information 240
 - starting in the service processor 241
 - stopping and exiting 244

- CCM support 29
- CCM User's Guide 30
- CCU-A 162
- CCU-A attachment 6
- CCU-B 162
- CCU-B attachment 6
- CD-ROM 39
- CD-ROM support 29
- CDLC 160, 163
- CF3745 29
- Channelized T1/E1 26
- Channellized T1/E1 support 27
- chargeable features 21
- CHPID 169
- Cipher Block Chaining (CBC) 25
- CIR 22, 37, 178, 183, 190
- CLA 152
- client/server 44
- CLLM 25
- CLP 17, 152
- committed burst size 178
- Committed Information Rate (CIR) 37, 178, 180
- committed rate measurement interval 178
- Communication Line Adapters (CLAs) 35
- communication line processor (CLP) 152
- communication rate (CR) 180
- communications rate (CR) 35
- community name 327, 328
- compression 194
- COMRATE 176, 178, 190, 191
- Config 220, 224, 258, 264
 - commands 331
- Configurable Quality of Service (QoS) 24
- configuration 237
 - activate 250
 - copy 249
 - create 249
 - create in service processor 250
 - service processor CCM version 250
 - create in Stand-alone 251
 - stand-alone CCM version 251
 - creation in different environments 250
 - export/import 250
 - modify 249
- configuring
 - ARP 338
 - BGP 336
 - BRS 339
 - controller 250
 - ESCON 341
 - frame relay 342
 - IP 332
 - OSPF 334
 - PPP 344
 - token-ring 341
 - with CCM 247

- connection network support for SVCs 24
- connectivity
 - adapter 19
 - maximum 18
 - network node 21
- controller bus and service processor (CBSP) 152
- controller configuration and management (CCM) 216, 237
- controller configuring 250
- controller expansion 16
- copy
 - configuration 249
- counting to infinity 95, 96
- coupler icons 245, 246
- couplers 152
- create a configuration 249

D

- D22510D 42
- D22510I 42
- D22510J 42
- D22510K 41
- D22560A 41
- D22560D 41
- D46100A 40, 173
- D46120 40
- D46120A 40
- D46130B 39
- Data Communications Equipment 34
- Data Encryption Standard (DES) 25
- data link control 48
- data link control identifier (DLCI) 191
- data link layer protocols 54
- Data Terminal Equipment 34
- database description 110
- database exchange 110
- DATABLK 176, 183
- DATABUFFERPOOLSIZE 236
- datagram 49
- DB2 26
- DCAF 216
- DCLI number 185
- default password 260
- default route 149
- DELAY 167
- designated router 105, 109
- despatching priority 235
- destination SAPs 173
- device number 167
- Dial-in support for SDLC PU Type 2 devices 27
- Dial-on-Demand 27
- different environments, configuration creation 250
- discard eligible (DE) 178
- discard frames 179

- distance vector routing 93
- distributed agent 227
- DLC parameters 167
- DLCI 178, 191
- DLCI range 36
- DLSw 25
- DLUR 28
- DoD (Dial-on-Demand) 27
- domain name 70
- dotted decimal 48
- DSAP 185, 187
- DTE 199
- DTE address 202
- DTE-to-DTE connections 202
- Dual Attach Station (DAS) 26
- dual CCU Attachment 7
- Dynamic address resolution 192
- dynamic reconfiguration (DR) 36
- dynamic routing protocol 167, 191
- DYNWIND 180

E

- EC level
 - D22510D 42
 - D22510I 42
 - D22510J 42
 - D22510K 41
 - D22560A 41
 - D22560D 41
 - D46100A 40, 173
 - D46120 40
 - D46120A 40
 - D46130B 39
 - D46130I 39
 - F12380x 38
- EMIF 165
- enable multicast emulation 191
- enable protocol broadcast 191
- Encryption Control Protocol (ECP) 25
- enhanced dynamic windowing algorithm 37
- Enterprise Extender 26
- environment
 - operating 238
 - service processor 238
 - stand-alone 238
- ERP 185
- ESCA 152
- ESCON
 - CDLC Configuration 162
 - host link 304
 - IP configuration 167, 168, 302
 - link station 304
 - MultiPath Channel (MPC+) 26
 - port 304

- ESCP 152
- Ethernet 16, 175
- Ethernet support 26
- euro-isdn 6
- excess burst size 178
- excess information rate (EIR) 181
- exiting from CCM 244
- export/import a configuration 250
- exterior gateway protocols 91

F

- F-Enet 26
- FDDI 26
- FDDI support 26
- FECN 25, 180
- filters 86, 275
- flooding 98, 112
- FR DCE (Data Communications Equipment) 34
- FR DTE 34
- fragmentation 53, 80
- frame relay 25, 176
 - boundary access node (BAN) 36
 - committed information rate (CIR) 180
 - communication rate (CR) 180
 - configuration 190, 191
 - frame handler (FRFH) 34, 35, 186
 - INN 187
 - over token-ring 35
 - Switching Equipment (FRSE) 34
 - Terminal Equipment (FRTE) 34
 - traffic scheduling 176
- frame switch 36
- FRFH 22, 34, 35, 36, 180, 181, 185, 186
- FRFH support 35
- FRSE 34
- FRTE 34, 180, 181
- FTP 23, 44
- Functional Overview 3

G

- GWCON 220, 222, 258, 262
 - commands 329

H

- hardware description 5, 7
- hardware requirements 238
 - 3746-900 16
 - 3746-950 18
- header information 245
- Hello protocol 108
- High-Performance Data Transport (HPDT) 26
- High-Speed Scanner (HSS) 34

- host link addresses 169
- hot-pluggability 5
- HPR 173, 185
- HPR MLTG 22, 39
- HPR SAP 173
- HPR-ERP 182
- HPR-ERP 182
- HSS 34
- HSSI 26
- HSSI support 26
- HTTP 23

I

- I/O addresses 162
- IARP table 192
- IBM 3746 Model 900 Expansion Unit 2
- IBM 3746 Nways Multinetwork Controller 2
- IBM 8229 16
- ICMP 65
- ICMP redirect
- icons, coupler 245, 246
- IGMP 115
- implicit focal point 25
- information area 246
- INN 185
- INN FRTE 34
- Inoperative error count 37
- installing
 - CCM 239
 - CCM in service processor environment 241
 - CCM in stand-alone environment 241
- interior gateway protocols 91
- Internal IP coupling 22
- Internet 43
- Internet protocol control protocol (IPCP) 194
- internetworking layer 47
- interoperability
 - BGP 149
 - OSPF 148
 - RIP 147
 - routing protocols 146
 - static routes 149
- intra-area router 104
- inverse ARP 177, 192
- IOCP 167, 306
- IP 162, 191, 195
 - address 71, 167
 - configuration 283, 293, 302
 - filter 312
 - FRTE 180
 - NCP 8
 - over frame relay 25, 36
 - parameters 307
 - routing 8, 58
 - routing feature 18

IP (*continued*)
 routing table 61
 traffic 187, 189

IP routing
 3746 131

IPCP 194

ISDN 6, 25

ISDN I.430 25

ISDN I.431 25

ITU-T I.370 179

L

LAN emulation 24

LAPB Parameters 204

LCP 194

LIC11 35

LIC12 35

LIC16 8, 22

LIC284 24

LIC286 26

LIC288 26

LIC289 26

LIC293 24

LIC294 24

LIC295 24

Link Access Procedure Balanced (LAP-B) 199

link control protocol (LCP) 194

link station 167

Link stations 162, 167

link-state acknowledgement 112

link-state advertisement 106

link-state propagation 111

link-state request 111

link-state routing 98

link-state update 111

LMI support 34

Local DTE Address 205, 210

locally administered MAC address 25

logical channel numbers (LCNs) 202, 205, 210

LPARs 165

M

main window

 becoming familiar with 244

 CCM 246

 menu choices 246

 working in 246

management password 233

maximum connectivity 18

maximum optional adapters 5

menu

 choices in main window 246

 pull-down 247

MIB 46, 225

minimum acceptable throughput 178

minimum information rate (MIR) 181

MMF 26

modify a configuration 249

monitoring

 ARP 339

 BGP 337

 BRS 340

 ESCON 341

 frame relay 343

 IP 333

 OSPF 335

 PPP 344

 token-ring 342

MOSPF 115

MOSS-E 214

MPC+ 26

MTU 206, 211, 236

multicast 115, 195, 314

multicast emulation 192

multicast group 115

MultiPath Channel (MPC) 26

MultiPath Channel (MPC+) 26

multiprotocol network backbone 185

MVS 306

N

Native APPN/HPR over ATM 24

Native HPR over ATM 23

NCP 34, 35, 195, 200

 V6R1 34

 V6R2 34

 V7R1 35

 V7R2 35

NCP V7R4 200

NCP V7R5 185

NCP V7R6 39

NCP/Boundary 185

neighbor interface state 108

neighbor router 104

neighbors 200

NetBIOS session alive spoofing 25

NetDispatcher 23

network control protocol (NCP) 195

Network Dependent 202

Network Information Center 92

network management 45

network node 2

network node base 16

network node connectivity 21

network node processor 3, 16, 151

next hop awareness 191

Next Hop Resolution Protocol (NHRP) 24

- next-hop address 163
- NHRP 24
- NLPID=X'08' 185
- NLPID=X'80' 185
- NLPID=X'CC' 185, 186
- NLPIDs 187
- NNP 16, 200
- non-ERP 185
- not NPSI 22
- NPM support 22
- NPSI 22, 200
- Numbering Plan Identification (NPI) 202

O

- OD media 30
- one-hop routing 24
- OPCON 219, 221, 258, 259
 - access failure 260
 - command level 258
 - commands 329
 - navigation 261
- Open Shortest Path First (OSPF) 102
- operating environments 238
- operation 237
- originate policy 121
- orphan circuit 177, 192
- OS/390 UNIX System Services UDP interface 26
- OS/390 V2R4 26
- OS/390 V2R5 26
- OSPF 200
 - 3746 IP 118
 - area 316
 - configuration 136, 313
 - multicast 314
 - route aggregation 317
 - virtual links 318
- OSPF protocol description 106
- OSPF terminology 102
- outboard frame switching 35
- Overview 1

P

- packet 49
- Packet Layer Protocol (PLP) 199
- packet switched data network (PSDN) 199
- partitioned mode 165
- password authentication 194
- Password Modes 233
- passwords 232
- path attributes 126
- PING 67, 192, 267
- PLP Modulo 205, 211
- PLP Parameters 205

- PLP Role 205, 210
- point-to-point protocol (PPP) 193
- poison reverse 101, 133
- port number 51, 64
- PPP 193
 - configuration 196, 197
- PPP Bandwidth Allocation Protocol/Bandwidth Allocation Control Protocol (BAP/BACP) 25
- PRI Euro-ISDN 22
- priority level 195
- protocol broadcast feature 192
- protocol stack 167
- proxy-arp 82
 - configuration 325
- pull-down menus 247
- PVC 202
- PVCs 34, 185

Q

- QoS 24

R

- read access 327
- receive policy 122
- remote DTE address 207, 212
- remote IP address 213
- remote login 44
- remote service facility (RSF) 216
- requirements, hardware and software 238
- restart 280
- RFC 1052 193
- RFC 1356 200
- RFC 1490 36
- RFC 1576 28
- RFC 1646 28
- RFC 1647 28
- RFC1490 186
- RFC1490 routed frame format 35
- RIP 200
 - 3746 IP 101
 - configuration 132, 318
- RIP outage-only advertisements 27
- RIP protocol description 99
- RIP V2 23
- RIP Version 2 23
- Rlogin 195
- Role 204
- route
 - default 61, 154
 - direct 61, 154
 - host 154
 - indirect 61
 - network 154

- route aggregation 79, 127, 317
- route caching 153
- ROUTED 99
- routed frame format 185, 186
- router ID 308
- router interface state 108
- router_id 68
- routing algorithms 92
- routing cache 153
- Routing caches 159
- Routing Information Protocol (RIP) 99
- routing table
 - display 159, 268

S

- SAP 26
- send policy 121
- sequence number 51
- Server Cache Synchronization Protocol (SCSP) 24
- service LAN 151, 283, 285
- service processor 3, 28, 151
 - environment 238
 - installing CCM in 241
 - starting CCM 241
- side cover 16
- Single Attach Station (SAS) 26
- SNA 182
- SNA peripheral device 35
- SNMP 45, 195, 215, 225, 327
 - configuration 326
- SNMP relay 215, 226
- SNMP trap 25, 227, 326
- socket 56
- software requirements 15, 238
 - 3746 IP 15
- split horizon 95, 101, 133
- split horizon with poison reverse 97
- SSAP=X'04' 187
- SSL 23
- stand-alone environment 238
 - installing CCM in 241
- starting CCM
 - in stand-alone CCM version 243
- static routes 93, 145, 167, 271, 310
- stopping and exiting from CCM 244
- subnet 74, 163
- subnet mask 74, 167
- supernetting 79
- SVC 202, 207, 212
- Switch-to-Switch Protocol 25
- TALK 6 261
- TCP 25
- TCP window sizes 236
- TCP/IP 43
- TCP/IP for MVS 162, 306
- TCP/IP stack 151
- TELNET 44, 195, 258
 - clients 215
 - command level 219
 - Config 220, 224
 - configuration 258
 - GWCON 220, 222
 - line command 329
 - line commands 218
 - OPCON 219, 221
 - operation 258
 - server 215, 259
- TIC1 36
- TIC2 36
- time to live 53
- TIMEOUT 167
- TME10 Remote Control 216
- TN3270 clients 28
- TN3270 gateway 28
- TOA/NPI 207, 212
- token ring 172
- token ring processor (TRP) 152
- token-ring
 - IP configuration 174, 293
- token-ring processor type 2 17
- TPF 162
- TRA 152
- traceroute 67, 270
- transport layer 47
- Trap 227, 326
- triggered updates 97, 101
- TRP 152
- TTL 53
- Type of Address (TOA) 202

U

- UDP 25, 64
- unit address 167, 170
- User Facilities 213
- user interface 244

V

- Van Jacobson 194
- variable subnet 75
- VCs 27
- viewing
 - product information for CCM 240
- virtual link 318

T

- TALK 5 261

VTAM 162

W

what you can do with a CCM configuration 249
working in the main window 246

X

X.121 202
X.25 22, 27
X.25 ODLC 200

Z

zero-hop routing 24

ITSO Redbook Evaluation

3746 Nways Controller Models 900 and 950: IP Implementation Guide
SG24-4845-01

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at <http://www.redbooks.ibm.com>
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

